

Europäisches Patentamt

(19)

European Patent Office

Office européen des brevets



(11)

EP 0 762 705 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:

12.03.1997 Bulletin 1997/11

(51) Int. Cl.⁶: H04L 29/06

(21) Application number: 96114234.6

(22) Date of filing: 05.09.1996

(84) Designated Contracting States:

AT BE CH DE DK ES FI FR GB GR IE IT LI LU MC
NL PT SE

(30) Priority: 08.09.1995 JP 231023/95

(71) Applicant: HITACHI, LTD.

Chiyoda-ku, Tokyo 100 (JP)

(72) Inventors:

- Takamoto, Yoshifumi
- Fuchu-Shi, Tokyo (JP)

• Kanai, Hiroki

Higashi-Yamato-Shi, Tokyo (JP)

• Takase, Tadahiro

Ohta-Ken, Tokyo (JP)

• Kitai, Katsuyoshi

Kokubunji-Shi, Tokyo (JP)

• Masuoka, Yoshimasa

Kodaira-Shi, Tokyo (JP)

(74) Representative: Beetz & Partner

Patentanwälte


Steinsdorfstrasse 10

80538 München (DE)

(54) Method for transmitting data via a network

(57) A communication controller (205) at the sending computer (201) divides data transferred from a host (101) into sub-ACK unit packets, and transfers them sequentially to a destination without waiting for the sub-ACK's being subsequently provided by the destination. A communication controller (206) at the receiving computer (202) issues a sub-ACK to the sending computer (201) for each of the sub-ACK unit packet, if the sub-ACK packet has been normally received and otherwise issues retransmission request for the sub-ACK unit packets and merges data included in the sub-ACK unit packets into the initial data, after they are normally received.

FIG. 2



EP 0 762 705 A2

Description

BACKGROUND OF THE INVENTION

The present invention relates to a method for transmitting data at high speed via a network and, more particularly, to a method for transmitting large quantities of data over long distances via a network.

Where computers can transmit data between them, they enjoy the advantage of sharing data. A plurality of computers may share one or more files therebetween by way of the file server. These functions are all implemented through the availability of data transmission between computers or between any one of the configured computers and a file server.

Data is transmitted generally via a network connected to a communication controller included in a computer. The communication controller is usually furnished on a single board. In operation, the communication controller receives a data transmission request or a data reception request from a computer and exchanges data with a communication target in units called packets. Each packet, which is the unit of data communication, contains user data and control information including destination and other information. The communication controller is incorporated in the host and is directly connected to an internal bus thereof. The host has a program called the operating system working therein. The operating system includes communication control software that controls the communication controller.

A typical network is what is known as the Fibre Channel. As discussed in "Fibre Channel Connection to the Future" (Fibre Channel Association, 1994, pp. 11-19), the Fibre Channel is available for communication between computers as well as between a computer and disk drives. The Fibre Channel is controlled functionally by layers FC0 through FC4 for each function of the fiber channel. Of these layers, FC0 defines the electrical characteristics of physical cables, FC1 stipulates the methods for encoding and decoding data, and FC2 specifies the data format relevant to the Fibre Channel. FC3 and FC4 designate logical communication methods. The Fibre Channel can connect a plurality of devices of different types, the capability being implemented by the layer FC4 converting data in format for use by a plurality of devices on the Fibre Channel. Illustratively, the Fibre Channel may be connected to disk drives as well as network devices. FC3 is the layer that provides common control over data communication. FC0 through FC4 are generally implemented by hardware, specifically on a board comprising a plurality of LSI's. With the Fibre Channel, the above-mentioned communication control software is a program that controls disk drives and networks. In operation, the communication control program issues requests to FC4 of the Fibre Channel. Upon receipt of a request from the communication control program, the Fibre Channel converts the data to its own data format in accordance with the FC4 provisions, forwards the data after conversion

through the layers FC3, FC2, FC1 and FC0, in that order, and sends to the destination device. Apart from the packets mentioned earlier, the Fibre Channel is capable of transferring data up to 2 KB long each as stipulated by FC2. On the Fibre Channel, any data exceeding the 2 KB length limit is divided into units of 2 KB each when transferred by the layer FC2.

The communication control software may be functionally divided into two major layers: a protocol processing program and a device control program. The protocol processing program receives data communication requests from a user program. For data transmission, the protocol processing program converts user data into packets according to the communication protocol specific to the network. For data reception, the protocol processing program extracts the user data from the transmitted packets. The device control program is a program that requests the communication controller to transmit packets. The primary function of the device control program is to control the hardware of the communication controller. In transmitting data, the device control program typically executes the following instruction string:

- (1) get_register (STATUS)
- (2) set_register (buf_addr)
- (3) set_register (pkt_length)
- (4) set_register (IN, DMA_GO)

The instruction (1) above reads the status of the communication controller and checks to see if the status is normal. The instruction (2) transmits the location of target data in the host to the communication controller. The instruction (3) notifies the communication controller of the length of the packet to be transmitted. The instruction (4) starts the hardware for transmitting the data from the host to the communication controller. The processing above allows the communication controller to analyze control information held in the packet and to transmit the target data to the destination of communication. Thereafter, the protocol processing program waits for an acknowledgment (abbreviated to ACK hereunder) to arrive from the destination indicating that the data has been normally transmitted. The above processes (1) through (4) are carried out by writing and reading information to and from control registers in the communication controller. In like manner, the device control program typically executes the following instruction string upon receiving data:

- (1) get_register (STATUS)
- (2) set_register (buf_addr)
- (3) set_register (pkt_length)
- (4) set_register (OUT, DMA_GO)

The instruction (1) above reads the status of the communication controller and checks to see if the status is normal. The instruction (2) transmits the location of target data in the host to the communication controller.

The instruction (3) notifies the communication controller of the length of the packet to be transmitted. The instruction (4) starts the hardware for transmitting the data from the communication controller to the host. Thereafter, the protocol processing program generates an ACK and sends it to the destination of communication indicating that the data has been normally transmitted. The processing above allows the communication controller to analyze the control information held in the packet and to transmit the target data to the destination of communication.

What characterizes the above type of communication is that an ACK is transmitted and received to verify the normal transfer of data every time a packet has been transmitted. With commonly implemented networks, the packet size is limited to be 1 to 4 kilobytes. To transmit data exceeding the size of one packet requires dividing the data into a plurality of packets that may be transmitted. In such a case, a plurality of ACK's are to be transmitted and received regarding each set of user data divided into a plurality of packets. After one packet is transmitted, the next packet cannot be transmitted unless and until the ACK for the most recently transmitted packet has been returned. This prohibition of transmission of the next packet is for communication. In data communication, if any packet is not transmitted normally, the same packet must be retransmitted. This requires retaining the most recently transmitted data in the memory on the transmitting side in preparation for possible retransmission. How long the data should remain in the memory is determined upon receipt of an ACK. That is, when the destination of communication has returned an ACK indicating the normal reception of the data, then the packet retained on the transmitting side is discarded.

One disadvantage of the above method of data transmission is that it takes time to transmit data. This is because the next packet cannot be transmitted until the ACK for the current packet is returned. A solution to this problem is described illustratively by Samuel J. Leffer et al., in "The Design and Implementation of the 4.3BSD UNIX Operating System" (Addison-Wesley Publishing Company, Inc., 1989, pp. 368-375). The solution involves transmitting a plurality of packets sequentially. In the above-cited reference, the protocol processing program stipulates that an ACK is to be issued for every n packets (i.e., n times the packet size). Within that unit of data transmission, packets may be transmitted and received sequentially. For example, suppose that the packet size is 1 KB, that an ACK is issued in units of 10 KB, and that 20 KB of user data are to be transmitted. In this case, the protocol processing program on the data transmitting side initially receives 10 KB of data from the user program. The protocol processing program divides the 10-KB data into 10 packets of 1 KB each. The 10 packets are sent consecutively to the destination. Having transmitted the 10 packets, the transmitting side waits for an ACK to be returned by the receiving side. The receiving side puts together the packets of 1 KB

each as they are continuously sent in. When the last 1-KB packet has arrived, the receiving side returns an ACK to the transmitting side. Upon receipt of the ACK, the transmitting side verifies that all data have been transmitted normally. If everything is normal, the protocol processing program of the transmitting side will perform the same process on the next 10-KB data. If the 10 packets were not transmitted normally, the same 10 packets are retransmitted. This method offers better data transfer performance than the method described earlier because packets can be transmitted consecutively within, say, the 10-packet range.

SUMMARY OF THE INVENTION

When a set of data is transmitted conventionally, the data is divided into packets and an ACK is returned every time a packet has been transmitted and received. If no ACK comes back, the next packet cannot be transmitted. The time required to return an ACK depends largely on the distance over which the data is transmitted. The longer the transmitting distance, the more prolonged the time required to return an ACK. This makes it difficult to transmit one packet after another continuously because an appreciable wait time can occur before an ACK arrives following the transmission of each packet. Where longer transmitting distances are involved, the network usability drops precipitously because of the prolonged ACK wait time. This impedes the network from being utilized effectively. The problem becomes more pronounced with high-speed data transmission media such as optical fiber cables. That is, the ability of the media to transmit data rapidly is canceled out by the increased time in which ACK's are awaited.

The solution to the problem proposed by the above-cited reference increases the size of the data for which one ACK is returned so that data may be transmitted at the high possible speed. More specifically, data items each smaller than an enlarged ACK are transmitted consecutively. By enlarging the ACK, the proposed method appears to provide faster data transmission than the earlier method. The greater the ACK, the more efficient data transmission is supposed to become. However, the proposed method is effective only if all data is transmitted normally. Long-distance communication is known to diminish the rate at which data is transmitted normally to the destination, mainly due to cable noise and data losses through network relay computers. Most network relay computers are connected to a plurality of networks. In such setups, memory constraints in any relay computer can truncate received data. The longer the transmitting distance is, the growing number of relay computers are installed to repeat data within or between networks. Numerous relay computers are liable to a growing probability of data loss. According to the method proposed by the above-cited reference, if any data item within a given ACK unit is lost, the entire ACK unit must be retransmitted. If retransmission overhead is taken into consideration in

an environment of a high probability of data loss, the performance of data transfer can be lower than that of the earlier method.

According to the present invention, one ACK unit is divided into a plurality of sub-ACK unit packets each furnished with a tag indicating the pre-division position. The divided packets are transmitted sequentially. Upon receipt of each divided packet, the receiving side communication controller returns a sub-ACK comprising the result of the reception to the transmitting side communication controller. Having received all packets, the receiving side communication controller determines the order of the packets as per the contents of their tags indicating the pre-division position and merges their data parts. The transmitting side checks the received sub-ACK's for the result of the reception to see if any packet need retransmission. Any packet is retransmitted if determined to be necessary. When all sub-ACK's have been normally received, the transmitting side puts them into a single ACK and notifies the request source thereof.

Because one ACK unit is divided into a plurality of sub-ACK unit packets for consecutive transmission, data transfer performance is enhanced. Since a packet retransmission request may be issued for each sub-ACK, drops in the transmitting performance are minimized in long-distance communication setups.

Other objects, features and advantages of the present invention will become apparent in the following specification and accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is an overall block diagram outlining a computer system which realizes a method for transmitting data via a network according to the present invention;

Fig. 2 is a method diagram for explaining the data transmission according to the present invention;

Fig. 3 is another diagram explaining the effects of the present invention;

Fig. 4 is a timing chart of the first embodiment;

Fig. 5 is a timing chart of one conventional method for transmitting data via a network;

Fig. 6 is a timing chart of another conventional method for transmitting data via a network;

Fig. 7 is a flowchart of steps showing how a user program of the first embodiment typically works;

Fig. 8 is a block diagram showing the constitution of the operating system of the first embodiment;

Fig. 9 is a flowchart of steps showing how a consecutive transmission controlling program of the first embodiment typically works;

Fig. 10 is another flowchart of steps also showing how the consecutive transmission controlling program typically works;

Fig. 11 is a block diagram depicting the transfer of data from a host to a communication controller;

Fig. 12 is a flowchart of steps showing how a packet

transmitting program of the first embodiment typically works;

Fig. 13 is a schematic view illustrating the structure of a transmission destination table;

Fig. 14 is a flowchart of steps showing how a packet dividing program of the first embodiment typically works;

Fig. 15 is a conceptual view depicting the transfer of data from the communication controller to a network;

Fig. 16 is a conceptual view indicating the constitution of a sub-ACK packet;

Fig. 17 is a flowchart of steps showing how a sub-ACK is received;

Fig. 18 is a conceptual view depicting the constitution of an ACK table;

Fig. 19 is a conceptual view outlining the effects of the first embodiment;

Fig. 20 is a flowchart of steps showing how a packet receiving program of the first embodiment typically works;

Fig. 21 is a conceptual view indicating how tags are removed and data are merged;

Fig. 22 is a conceptual view sketching how the invention is applied to a disk drive arrangement;

Fig. 23 is an overall block diagram of a second embodiment of the invention;

Fig. 24 is a conceptual view portraying the constitution of a packet for use with the second embodiment;

Fig. 25 is a flowchart of steps showing how a divided packet transmitting program of the second embodiment typically works;

Fig. 26 is a flowchart of steps showing how a divided packet receiving program of the second embodiment typically works; and

Fig. 27 is a conceptual view indicating how packets are consecutively transmitted in a multiple port setup of the second embodiment.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

First Embodiment

Preferred embodiments of the invention will now be described in detail with reference to the accompanying drawings.

Fig. 1 is an overall block diagram outlining a computer system which adopts a method for transmitting data via a network according to the present invention. In Fig. 1, reference numeral 101 is a host and 105 is a communication controller. The host 101 comprises a processor 102 for executing user programs and so on, a host memory 103 for accommodating the user program, user data, etc., and a bus 104. The bus 104 is connected to a communication controller 105 or the like for communicating with the processor 102 host memory 103 and outside of the computer. The communication controller

105 includes a bus interface 106 that communicates with the bus inside the host 101, a control processor 107 that controls the hardware in the communication controller 105, a control memory 108 that stores control software to be executed by the control processor 107, a communication interface 119 that communicate with and a data memory 121 and externally connected devices, a communication load measuring device 120 that measures the load of the network to which the communication controller 105 is connected, the data memory 121, a data bus 117 for data transmission, a control bus 118 for transmitting control information, a backup device 122 that backs up the contents of the data memory 121, a transmitter 126, and a receiver 125. The transmitter 126 and receiver 125 are connected to an external network. The bus interface 106 comprises a buffer 115 that temporarily retains data transferred between the host memory 103 and the communication controller 105, and a DMAC 116 that transfers data between the host memory 103 and the communication controller 105. The communication interface 119 is made up of a buffer 124 that temporarily holds data transferred between the data memory 121 and an externally connected device, and a DMAC 123 that transfers data between the data memory 121 and the external device. The control memory 108 stores control software to be executed by the control processor 107. Reference numerals 109 and 112 denote control software dealing with packet transmission and packet reception, respectively. The packet transmitting program 109 comprises a packet dividing program 110 and a tag adding program 111, etc. The packet receiving program 112 includes a packet merging program 113 and a tag removing program 114, etc.

How the system of Fig. 1 works will now be described by referring to Fig. 2. What is shown in Fig. 2 is a process of data transmission and reception between a host 201 (corresponding to the host 101 in Fig. 1) and a host 202 (corresponding to the host 101 in Fig. 1), both connected to a network 216. Specifically, the host 201 transmits data 203 constituting an ACK unit of the host 201, the data 203 being received (as data 204) by the computer 202. The ACK unit is a unit of data for which an ACK must be returned before the next packet of data can be transmitted. The computers 201 and 202 are connected to communication controllers 205 and 206 respectively. The data 203 sent from the computer 201 is divided in the communication controller 205 into three data 209 each constituting a sub-ACK unit. Each divided data 209 is constructed as a packet 208.

This setup may be elaborated in connection with the Fibre Channel as follows. The Fibre Channel is connected to disk drives and network devices. An example of data communication between computers in the Fibre Channel will be explained by reference to Fig. 2. The control programs to control the connected devices run on the computers 201 and 202. Each of the control programs to control the disk drives and networks issues a

request for data (corresponding to 203) to the communication controller (corresponding to 205) of the Fibre Channel. In response, the communication controller (205) divides the data (203) into units of up to 2 KB each, the maximum unit of data transfer by the Fibre Channel, to be sent as divided data (209). The divided data 209 thus transferred are received by the destination communication controller (corresponding to 206) of the Fibre Channel and reconstructed therein as packets 208 for transfer to the computer 202. Specifically, of the Fibre Channel layers FC0 through FC4, the layer FC4 takes care of data division control. These arrangements make it possible to practice the invention without modifying the existing control software.

Returning to Fig. 2, besides its user data (i.e., divided data) 209, each packet 208 is furnished with tags such as a destination ID 210, a flag 211 indicating a division of the data 203, a packet ID 212, a division count 213 (number of divided data), a division ID 214 indicating where the divided data is located in the post-division data order, and a transmission time 215 at which the packet is transmitted. Upon reception of each data, the communication controller 206 removes the tags from each packet so as to reconstruct the initial data 204. During data reconstruction, the communication controller 206 refers to the division flag 211, packet ID 212, division count 213 and division ID 214 indicating the post-division data order. For each packet, the communication controller 206 returns a sub-ACK 207 indicating completion of data transmission. If any one of the three transmitted packets has developed a transmission error, the occurrence of the error is written into the corresponding sub-ACK 207. In response, the communication controller 205 retransmits only the packet corresponding to the sub-ACK in question. Where any of the divided packets has developed transmission error, only the faulty packet is retransmitted; there is no need to retransmit all packets 208. When all sub-ACK's are normally returned, the three sub-ACK's are integrated into one ACK and the host 101 is notified of the normal completion of data transmission. Because the target data is divided into sub-ACK units each furnished with a division flag 211, a packet ID 212, a division count 213 and a division ID 214 indicating the post-division data order, the divided data are transmitted and received as an independent packet each through parallel packet transmission. The sub-ACK's are merged into one ACK which is then transferred to the source of the request. The procedure here is equivalent to conventional communication protocol and contributes to minimizing the amount of changes needed in the software of the host where the invention is applied. Details of the hardware and software of the above setup will be described later.

Fig. 3 is a diagram for showing effects of the present invention. As with the setup of Fig. 2, the example of Fig. 3 divides an ACK unit of user data into three sub-ACK units (319) for consecutive packet transmission. Specifically, a communication controller 302 trans-

fers user data 301 to be received by another communication controller 317. The communication controller 302 divides the data 301 into three packets each furnished with tags. The packets are outputted consecutively onto a transmission path 306 (308 through 310). Packets 311 through 313 denote packets which include divided data of user data that were outputted earlier than the user data 301. The communication controller 317 removes tags from the packets 311 to 313 which include the previously outputted user data to reconstruct the initial data user (318). For the packets 311 through 313, the communication controller 317 returns sub-ACK's 314 through 316 respectively. When all sub-ACK units have been returned, the sub-ACK's are merged into one ACK (320) which is sent to the source of the request. Unlike the conventional arrangement in which the next packet cannot be transmitted until an ACK is returned for the current packet, the embodiment of the invention waits for the return of an ACK only between the packets 308 through 310 on the one hand, and the packets 311 through 313 on the other. The result is data transmission at an enhanced speed. Because the sub-ACK's are returned over a path different from the transmission path, there is no possibility of collision between the data and sub-ACK's. This is another advantage of the invention.

Fig. 4 is a timing chart showing the flow of control of the first embodiment. Fig. 4 depicts an example in which a host 401 transfers data 405 to another host 405. The user data 405 sent from the host 401 to a communication controller 402 is divided by the latter into a plurality of sub-ACK units (408 through 411) which are transmitted consecutively to the destination. Upon receiving the data, a communication controller 403 returns a sub-ACK for each of the divided sub-ACK unit packets (412 through 414). The communication controller 403 further merges the divided packets for data transfer to the host (407). Meanwhile, the communication controller 402 receives a plurality of sub-ACK's and merges them into one ACK which is returned to the host. In this manner, the user data is divided and transmitted in units of sub-ACK's, each data division being transmittable independently, whereby the data transfer performance is improved.

Below is a comparison of the first embodiment with typical prior art communication methods whose workings are illustrated in Figs. 5 and 6. Fig. 5 shows a conventional method wherein an ACK is returned for each of the packets (505, 507) transmitted by a host 501. According to the method of Fig. 5, hosts 501 and 504 or communication controllers 502 and 503 transmit data and ACK for each data transfer request from a user program. Because it takes time before each ACK is returned with this method, a prolonged time period is required before a succeeding packet is transmitted. Fig. 6 depicts another conventional method which improves transfer performance by dividing a set of user data into a plurality of packets and by transmitting those packets. According to the method of Fig. 6, a data receiving com-

munication control board 603 returns only one ACK for the entire set of user data (605 or 607). This means that if any one of the divided packets has developed transmission error, all packets need to be retransmitted. Where long communication distances are involved, the conventional method is subject to a growing probability of data losses and can diminish the performance of data transfer instead of improving it. These pitfalls are circumvented by the first embodiment of the invention using sub-ACK units. The sub-ACK units make it possible to verify which of the divided packets has developed an error. Therefore, even transmission data over long distances requires small amount of data to be retransmitted and thus enjoy higher speeds of data transfer.

How the first embodiment works will now be described in more detail by referring to the accompanying drawings from Fig. 7 on. Fig. 7 shows an example of user programs. This is an example of data communication between the user program on the transmitting side (701 through 703) and the user program on the receiving side (704 through 706). Steps 701 and 704 define the environment necessary for communication. At this point, communication is yet to take place between the two programs. Steps 702 and 705 check to see if a communication target and is ready to communicate. Step 703 transmits user data to the communication target checked in step 702. There is no need for the user program to become aware of the packet size in effect. That is, there are no concept of a packet size or an ACK unit to be taken into account at this stage of communication. The packet size and ACK unit are concepts defined by the operating system 2403, to be described later.

Fig. 8 shows the constitution of the programs in a host memory 2401 (corresponding to the host memory 103 in Fig. 1). When a user program 2402 issues a communication request, a processing request is transferred to the operating system 2403. The communication processing is carried out by a packet transmitting program 2404 and a packet receiving program 2405 inside the operating system 2405. The packet transmitting program 2404 and packet receiving program 2405 are known programs that define ACK units. Whereas the conventional system has a communication controller 2408 (corresponding to the communication controller 105 of Fig. 1) directly controlled by the packet transmitting program 2404 and packet receiving program 2405, the first embodiment of the invention has its packet transmitting program 2404 connected to a consecutive transmission controlling program 2406 and a transmission incomplete data buffer 2407. The consecutive transmission controlling program 2406 is capable of enlarging the ACK unit of the packet transmitting program 2404. Details of the consecutive transmission controlling program 2406 and transmission incomplete data buffer 2407 will be described next.

Fig. 9 is a flowchart of the consecutive transmission controlling program 2406 of the host at data transmission. The consecutive transmission controlling program 2406 is called up by the packet transmitting program

2404 of Fig. 8. Step 2501, temporarily places in the transmission incomplete data buffer 2407 the data for which a transmission request has been issued by the packet transmitting program 2404. Step 2502 issues a transmission request to the communication controller 2408. Details of this step will be described later. The consecutive transmission controlling program 2406 notifies the packet transmitting program 2404 is notified of the completion of data transmission at this moment by the step 2503. This allows the packet transmitting program 2404 to regard as if the ACK has been returned for the transmitted data, and to request transmission of the next data. If more data to be transmitted is found to exist in step 2504, the process starting from the step 2501 is again repeated. The transmission incomplete data buffer 2407 retains its contents until the transmission is normally completed, in preparation for possible retransmission of lost packet data. Whereas the procedure of data retention is performed traditionally by the packet transmitting program 2404, the program 2404 here discards the data because the process is considered normally terminated by the step 2503.

Fig. 10 is a flowchart of the consecutive transmission controlling program 2406 at reception. In step 2601, an ACK is received from the transmitting side. In step 2602, a check is made to see if all user data has been transmitted. If the transmission is normal, the applicable data in the transmission incomplete data buffer 2407 is erased therefrom in step 2603. The processes of Figs. 9 and 10 allow the ACK unit to be enlarged without modifying the existing programs. This is another advantage of the present invention.

Fig. 11 is a block diagram depicting how data is transmitted from a host 801 (corresponding to the host 101 of Fig. 1) to a communication controller 802 (corresponding to the communication controller 105 of Fig. 1). Reference numeral 803 points to details of the processing of step 2502 carried out by the consecutive transmission controlling program 2406.

(1) get_register (STATUS)

In this processing, the status of the communication controller 802 is acquired. When the status is found to be normal, the transmission processing continues.

(2) set_register (buf_addr)

In this processing, the start address of the transfer target data in the host 801 is written into one of a group of control registers (807 through 810) in the communication controller 802.

(3) set_register (pkt_length)

In this processing, the length of the transfer data in the host 801 is written into one register of the control register group (807 through 810) in the communication controller 802.

(4) set_register (IN, DMA_GO)

In this processing, an instruction issued to the DMAC 116 of Fig. 1 for data transfer to a buffer 811. The transferred data has the start address and the length both held in the control registers. The data in the host 801 is in packet format and comprises control information 804 such as the destination ID in addition to the user data 805.

Fig. 12 is a flowchart of the packet transmitting program 109 of the communication controller 105 of Fig. 1. In step 901, the data placed into the buffer 811 by the processing shown in Fig. 11 is transferred to the data memory 121. Following step 901, the communication controller 107 can control the data. In step 902, the packet data is divided into sub-ACK unit packets by the packet transmitting program 109 in accordance with a division unit. How the division unit is determined will be described later with reference to Fig. 14. In step 903, each of the sub-ACK unit packets derived through division from the packet data is furnished with tags 209 through 215 in Fig. 2. In step 904, the sub-ACK unit packets are transmitted to the destination. In this step, the control processor 107 transfers the packets to a buffer 124 of Fig. 1 via a communication interface 119 of Fig. 1 by use of a data transfer device 123 of Fig. 1. Then, the transmitter 126 transmits the sub-ACK unit packets according to the relevant protocol. In step 905, a check is made to see if all sub-ACK unit packets have been transmitted. If any sub-ACK unit packet has yet to be transmitted, step 904 is reached again and the packet in question is retransmitted. In steps 906, 909 and 907, sub-ACK's are awaited for all divided packets transmitted. After these steps have verified the sub-ACK's for all normally transmitted packets in sub-ACK units, it becomes possible to merge the sub-ACK's. If any sub-ACK indicates an abnormally transmitted packet, only the packet in question is retransmitted in step 909. When the sub-ACK's have been received from the destination indicating that all divided packets have been normally transmitted, the sub-ACK's are merged. In step 908, the packet transmitting program 109 notifies the host of the completion of data transfer.

How a packet is divided in step 902 of Fig. 12 will now be described in more detail with reference to Figs. 13 and 14. Fig. 13 is a schematic view illustrating the structure of a transmission destination table which is held in the control memory 108, is controlled by the communication controller 105, and accommodates information for use upon packet division. In the table, column 1001 holds destination ID's representing previously accessed transmission destinations. New destinations may be added to the table. Column 1002 retains units of division. Different destinations have different distances to cover and need to be assigned an optimum division unit each. Column 1003 stores the response time registered when each destination was most recently accessed. The response time is the time required to send data of a specific length. Column 1004

holds the communication loads specific to the destinations. Column 1005 accommodates the direction of division for each destination. The direction of division indicates whether the current unit of division is to be enlarged or reduced. An indication "+1" means that the division unit is to be enlarged and "-1" means the division is to be reduced. Every time data is transferred, the response time column 1003 and communication load column 1004 are referenced and the division unit is updated in such a manner that the data transfer performance will be improved. Details of this process will be described later with reference to Fig. 14. The division unit is dynamically changed upon every access so that an optimum division unit will be gradually approximated. The transmission destination table is stored in the data memory 121 and backed up by the backup device 122. Thus the table is preserved unless and until the user forcibly resets the table contents. The use of the transmission destination table permits data transfer of constantly optimum performance.

Fig. 14 is a flowchart of the packet dividing step 902 of Fig. 12. In step 1101, the transmission destination table of Fig. 13 is referenced. Upon data transfer to the communication controller 105, the host 101 attaches a destination ID (804 in Fig. 11) to the user data. In step 1102, a check is made to see if the attached destination ID is already stored in the transmission destination table of Fig. 13. If the destination ID is not found in the table, step 1104 is reached in which the destination is added anew to the table. In step 1105, the division unit of the newly added destination is set to the minimum division unit previously defined by the system. Step 1103 is reached if the destination is found in the transmission destination table. In this step, the division unit corresponding to the destination ID is retrieved from the table. In other words, if the destination is already stored in the transmission destination table, the corresponding stored division unit is used unchanged. Otherwise the minimum division unit is set for the newly added destination.

Fig. 15 depicts detailed operations of the packet transmitting step 904 in Fig. 12. The operations are performed to transfer data from the communication controller 105 to an external network. The instructions involved are as follows.

(1) get_register (STATUS)

In this process, the status of the communication interface 119 is acquired. If the status is normal, the transmission processing continues.

(2) set_register (buf_addr)

In this process, the first address of the transfer data 1206 in the data memory 121 is written into one of a group of control registers 1207 in the communication interface 119.

(3) set_register (pkt_length)

In this process, the length of the transfer data 1206 in the data memory 121 is written into a register of the control register group 1207 in the communication interface 119.

(4) set_register (OUT, DMA_GO)

This process is performed by the instruction issued to the DMAC 123 of Fig. 1 for data transfer to the buffer 124. The transferred data has the same length as that held in the control registers from the first address on. The processes above transfer the data 1206 from the data memory 121 to the buffer 124. In response, the communication interface 119 passes the data on to the transmitter 126 for packet transmission to the network.

Fig. 16 indicates the constitution of a sub-ACK packet returned for a packet by the receiving side to the transmitting side. The sub-ACK comprises the result of transmission 1301, a destination ID 1302 representing the destination of transmission, a division flag 1303 indicating a division, a packet ID 1304 identifying the packet, a division count 1305 indicating the number of packet divisions, a division ID 1306 indicating the post-division sequence, and a transmission time 1307 at which the transmitting side transmitted the packet. The result of transmission 1301 shows whether or not the packet has been normally received. Depending on the transmission result, the transmitting side determines whether or not to retransmit the packet in question.

It should be noted that the destination ID 1302, division flag 1303, packet ID 1304, division count 1305, division ID 1306 and transmission time 1307 are the same as those attached to the data packet sent from the transmitting side. The attached information of the data packet is tagged unchanged to the sub-ACK returned for the data packet.

Fig. 17 is a flowchart of the sub-ACK processing step 909 in Fig. 12 for sub-ACK's of Fig. 16 received by the packet transmitting side. In step 1401, a check is made by referencing the transmission result 1301 of the sub-ACK to see if the divided packet has been normally transmitted. If the packet is found to be normally transmitted, step 1402 is reached; otherwise step 1406 is reached. In step 1406, a retransmission request is issued to the packet transmitting program of Fig. 12 (specifically step 904) regarding the abnormally transmitted packet. Step 904 of Fig. 12 started here is terminated when the packet has been retransmitted. At this point, a retransmission count is updated as will be described later by referring to Fig. 18 in connection with a sub-ACK control table. In step 1402, the time required for the transmission is calculated. The result of the calculation serves as a parameter for determining an optimum division unit. The calculation may be performed by finding the difference between the transmission time 1307 in Fig. 16 and the current time. Specific means for knowing the time may illustratively be a timer incorpo-

rated in the communication controller. In step 1403, the current network load is obtained by the communication load measuring device 120 of Fig. 1. The communication load measuring device 120 is designed to acquire the current load of the network to which the communication controller 105 is connected. The network load is detected by checking to see how many packets have been received by the receiver 125 within a unit time. The load thus acquired serves as a parameter for determining an optimum division unit. In step 1404, a check is made to see if the division unit 1002 currently held in the transmission destination table (Fig. 13) is appropriate. The check in step 1404 is made by comparing the response time 1003 and communication load 1004 in the transmission destination table (Fig. 13) on the one hand, with the currently acquired response time and communication load on the other hand. The division unit needs to be updated in two cases: when the current communication load is smaller than the communication load 1004 held in the transmission destination table (Fig. 13) and the current response time is smaller than the stored response time; or when the current communication load is smaller than the communication load 1004 in the transmission destination table and the currently obtained response time is larger than the stored response time. In the former case, the stored division unit is updated in the forward direction of the currently obtained division unit. In the latter case, the update is performed in the opposite direction of the currently acquired division unit. The updating is executed in step 1405. Specifically, where the update is in the forward direction (i.e., +), the stored division unit is enlarged. For example, the stored division unit of 2KB is replaced by 4KB. If the update is in the opposite direction (i.e., -), the stored division unit is reduced. Illustratively, the stored division unit of 4KB is replaced by 2KB. Such updates make it possible to reduce the division unit for long distances and to enlarge it for short distances. This is because the long-distance transmission takes more time for sub-ACK's to be returned and has a good response time with the smallest possible division unit.

Where the retransmission count, to be described with reference to Fig. 18, is used as one parameter for dividing the packet, the data transfer performance may be improved further. This is particularly effective where many errors are prone to occur during data communication with a specific host. When the retransmission count is found to increase, the retransmission time can be shortened by reducing the division unit. For example, if numerous retransmission requests occur with the division unit of 8KB in effect, the retransmission overhead may be lowered by setting a smaller division unit such as 4KB. How the retransmission process is carried out will be described later in detail with reference to Fig. 19.

Below is a description of the network load. The network is not for use by a single user; a plurality of users are expected to transfer data concurrently over the network. In such cases, the network load will rise causing the response of data transfer to worsen. The deteriora-

tion of response is not related to the distance of communication. Thus by detecting the network load and using it as a parameter for determining the division unit, the first embodiment of the invention can discriminate those variations in response which are not attributable to distance. Although the first embodiment has the communication load column 1004 established in the transmission destination table (Fig. 13), the column may be replaced alternatively by a column for holding communication distances. If implemented, the communication distance column will allow the division unit 1002 to be obtained more accurately. The communication distances are to be set previously by the user for the stored destinations, and new destinations may be added to the table and matched with the correspondingly furnished communication distances to ensure higher performance.

How a plurality of sub-ACK's are awaited will now be described with reference to Fig. 18. The process involves executing steps 906 and 907 of Fig. 12 to wait for the divided sub-ACK's to arrive. Which sub-ACK has been received at any given moment is known by use of a sub-ACK control table of Fig. 18 established in the data memory 121. Column 1501 accommodates packet ID's representing the same contents of the packet ID's 212 in the packets 208 of Fig. 2. Likewise, column 1502 holds division ID's that are the same as the division ID's 214 in the packets 208. Each division ID 1502 corresponds to a sub-ACK. Information on whether or not a sub-ACK for a specific division ID 1502 has been received is written to column 1503. If all sub-ACK's are found to be returned, that means the completion of data transfer, which verifies the end of step 906. The sub-ACK control table is generated upon division of a packet (step 902 in Fig. 12).

Fig. 19 shows that data may be retransmitted at high speed when the division unit is small. Reference numeral 1604 indicates a set of data transmitted in a large unit, and 1605 denotes a packet of the same data retransmitted. Reference numerals 1606 through 1609 indicate the data being transmitted in smaller units, and 1610 stands for one packet of the retransmitted data. In many cases where retransmission is requested, only a part of the data is erroneous. In such cases, the entire data must be retransmitted (1605) if the packet is large enough (1604). By contrast, if the division unit is small, it may take more time to transmit data but the subsequent retransmission thereof will end quickly because the latter process involves an appreciably low amount of data. As a result, the overall time for data transfer is shortened.

Fig. 20 is a flowchart of the packet receiving program 112 of the communication controller 105 of Fig. 1. In step 1701, a packet is received via the communication interface 119. In step 1702, a check is made to see if the packet has been normally received. If the reception of the packet is found to be normal, step 1703 is reached; otherwise step 1704 is reached. Whether or not the packet is normally received is verified illustratively by use of a scheme wherein the transmitting side

attaches parity or cyclic redundancy check bits to the transmitted data and the receiving side checks the attached bits. In step 1703, a sub-ACK indicating the normal reception of data is returned to the transmitting side. In step 1704, a sub-ACK indicating an abnormal receipt of data is returned to the transmitting side. In step 1705, a check is made to see if all sub-ACK units have been received. The check in step 1705 is accomplished by referring to the division flag 211, packet ID 212, division count 213 and division ID 214 shown in Fig. 2. In step 1706, the tags 210 through 215 in Fig. 2 are removed from the transmitted data so that only the data part is left. The removed tags are retained corresponding to the data and are used in the next step (1707). In step 1707, the divided data are merged according to the division ID 214. In step 1708, the merged data is transferred to the host and the packet receiving process is terminated.

Fig. 21 is a conceptual view indicating how divided data are merged. The example of Fig. 21 is one in which three packets 1802, 1803 and 1804 are merged. The divided data extracted from the three packets are merged into one set of data (1805). The initial data, i.e., the data transmitted by the host of the transmitting side to the communication controller, is reconstructed by merging the divided data in accordance with division ID's 1808.

Fig. 22 is a conceptual view sketching how the invention is applied illustratively to a disk drive arrangement. Whereas the description of the invention above has centered on data communication between hosts, the invention is intended to boost the speed of data communication and as such can be applied generally to data communication setups including the disk drive arrangement. Fig. 22 shows a process in which a computer 1901 (host) connected to a network 1909 transmits data 1902 (in the ACK unit for the host) to a disk controller 1905 of a disk drive 1904. Reference numeral 1906 denotes the data received by the disk controller 1905. The computer 1901 is connected to a communication controller 1903, and the disk controller 1905 is connected to a communication controller 1907. The data 1902 sent from the computer 1901 is divided by the communication controller 1903 into three sub-ACK unit packets 1910. As in the example of Fig. 2, each packet 1910 comprises divided user data and is furnished with tags such as the destination ID, division flag indicating a division, packet ID, division count, division ID indicating the post-division sequence, and transmission time at which the packet was transmitted. Upon receipt of the packets, the communication controller 1907 removes the tags from the packets to reconstruct the initial data 1906. For data reconstruction, the communication controller 1907 references the division flag, packet ID, division count, and division ID indicating the post-division sequence of each packet as in the case of Fig. 2. For each packet received, the communication controller 1907 returns a sub-ACK 1908 indicating the completion of packet transmission. If any one of the three transmit-

ted packets has developed transmission error, the occurrence of the error is written to the sub-ACK 1908 corresponding to the abnormally transmitted packet. In response to the sub-ACK, the communication controller 1903 retransmits only the packet corresponding to that sub-ACK. There is no need to retransmit all packets 1902 constituting the entire data where any one of the packets has developed transmission error. When the sub-ACK's are returned indicating the normal reception of all divided packets by the receiving side, the three sub-ACK's are merged into one ACK which notifies the host 1901 of the completion of data transmission. Because user data is divided into sub-ACK unit packets each provided with the division flag, packet ID, division count, and division ID indicating the post-division sequence ID, the packets are transmitted and received independently and hence in parallel. Where a plurality of disk drives 1904 are connected to the disk controller 1905, the disk controller 1905 may process a plurality of disk I/O requests concurrently. This is a significant advantage available with the embodiment of the invention allowing the host 1901 to transfer data at high speed. Because sub-ACK units are adopted and they are merged before transfer to the host computer, i.e., to the source of the request, the protocol requirements of the first embodiment are the same as those of conventional communication protocols and thus involve only a minimum of modifications in the software of the host upon implementation. In addition, the greater the amount of transfer data, the more efficient the parallel transfer of data carried out for transmission and reception with the first embodiment.

Second Embodiment

Fig. 23 is a block diagram of a software implementation of the invention. In Fig. 23, reference numeral 2001 is a host and 2011 is a communication controller. Working inside the host 2001 are a user program 2002 and the operating system 2003. The operating system 2003 includes a communication controlling program 2004 which in turn comprises a division controlling program 2012, a packet transmitting program 2005 and a packet receiving program 2006. The division controlling program 2012 incorporates a divided packet transmitting program 2007 and a divided packet receiving program 2009. A communication request issued by the user program 2002 to the network is first received by the communication controlling program 2004. Later, control is passed on to the division controlling program 2012 whose detailed workings will be described later. The packet transmitting program 2005 and packet receiving program 2006 are known programs, and the communication controller 2011 is a known device.

The second embodiment utilizes a conventional system when implemented. The packet transmitting program 2005 and packet receiving program 2006 deal with packets each comprising two major parts, a data part 2101 and a packet header part 2102, as shown in

Fig. 24. The data part 2101 is provided by the division controlling program 2012. Whereas the first embodiment attaches specific tags to the packet header part, the second embodiment cannot modify the packet header because it utilizes the existing system. Thus the second embodiment attaches to the data part 2101 those tags (2104 through 2109) that are necessary for dividing packets. Besides user data 2103, each packet is furnished with such tags as a division flag 2104 indicating a division, a division count 2105 representing the number of packet divisions, a division ID 2106 indicating the post-division sequence, and a transmission time 2107 at which the packet was transmitted.

Fig. 25 is a flowchart of steps showing how the divided packet transmitting program 2007 typically operates. In step 2201, the packet is divided into sub-ACK units. The details of step 2201 are the same as those of the first embodiment in Fig. 14. In step 2202, each of the divided sub-ACK unit packets is given the tags 2104 through 2107 shown in Fig. 24. In step 2203, the sub-ACK unit packets are transferred to the destination. In step 2204, a check is made to see if all sub-ACK units have been transmitted. If there exist any sub-ACK unit packets yet to be transmitted, step 2203 is reached again and executed. In steps 2205, 2208 and 2206, sub-ACK's are awaited for all transmitted packets. When sub-ACK's have been received indicating the normal reception of all the packets, it is now possible to merge the sub-ACK's. This allows the user program 2202 to remain unaware of the division of the packet. If there is a sub-ACK indicating an abnormally transmitted packet, that packet alone is retransmitted in step 2208. Step 2208 is the same in function as step 909 in Fig. 12. When all sub-ACK's have been merged, step 2207 is reached in which the user program 2002 is notified of the completion of data transfer.

Fig. 26 is a flowchart of steps of the divided packet receiving program 2009. In step 2301, a packet is received via the packet receiving program 2006. In step 2302, a check is made to see if the packet has been normally received. If the receipt of the packet is found to be normal, step 2303 is reached; otherwise step 2304 is reached. In step 2303, a sub-ACK indicating the normal reception of data is returned to the transmitting side. In step 2304, a sub-ACK indicating an abnormal receipt of data is sent back to the transmitting side. In step 2305, a check is made to see if all sub-ACK unit data have been received. The check in step 2305 is accomplished by referring to the division flag 2104, division count 2105 and division ID 2106 of each packet shown in Fig. 24. In step 2306, the tags 2104 through 2107 in Fig. 24 are removed from each packet so that only the data part is left. In step 2307, the divided data are merged in accordance with the division ID's 2106. In step 2308, the merged data is transferred to the user program 2002 and execution of the divided packet receiving program is 2009 terminated.

Fig. 27 illustrates schematically the principles of packet division of the second embodiment. The known

packet transmitting program 2005 and packet receiving program 2006 are each provided with a plurality of I/O ports for communication. The I/O ports are made to correspond with the destinations of communication on a one-to-one basis. Because the next packet cannot be transmitted through the same port until the normal receipt of the current packet is ascertained by a returned ACK, the packet dividing function of the invention cannot be implemented using the same port. Thus the second embodiment has at least two ports (2705, 2706, 2716, 2717) assigned to each of the communication controlling program 2701 of the transmitting side and the communication controlling program 2702 of the receiving side dealing with one set of user data. In consecutive packet transmission, this setup makes it possible to transfer the next packet independently of the ACK returned for the most recently transmitted packet. The division controlling programs 2704 and 2719 transfer user data 2703 and 2718 independently through each of the ports in use 2711 and 2712. It is also possible to return sub-ACK's 2713 and 2714 independently through each of the ports. With the second embodiment, communication controllers 2720 and 2719 need only perform conventional operations. Other detailed processes and the table arrangements are the same as those of the first embodiment. The second embodiment has the advantage of eliminating the need for any modifications in the conventional hardware or in the current operating system 2003.

As described, the invention minimizes deterioration of the data transfer performance where long distances are involved for data communication.

Claims

1. A method for transmitting data via a network wherein the transmission of a data packet by a transmitting side (205(Fig. 2); 1903(Fig.22); 2704(Fig.27)) is acknowledged by an acknowledgement packet returned by a receiving side (206(Fig. 2); 1907(Fig.22); 2719(Fig.27)) before the next packet can be transmitted, said transmitting side carrying out the steps of:

- (1) dividing each data packet into a plurality of data packets;
- (2) furnishing each of the divided data packets with an identifier indicating the pre-division position in sequence of the data packet in question; and
- (3) transmitting said divided data packets without waiting for an acknowledgement packet to arrive for each data packet transmitted;

said receiving side (206, 1907, 2719) carrying out the steps of:

 - (4) merging said divided data packets upon receipt in accordance with the identifiers thereof; and
 - (5) returning, for each of said divided data

packets received, an acknowledgement packet provided with receipt information indicating whether the data packet in question has been normally received.

2. A method for transmitting data via a network according to claim 1, wherein, after the normal transmission of all of said divided data packets has been verified, one acknowledgement packet is generated and returned to the source which requested the transfer of the data packet constituted by said divided data packets.
3. A method for transmitting data via a network according to claim 2, packet reception information included in said acknowledgment packet is referenced in order to retransmit only the packets not normally transmitted.
4. A method for transmitting data via a network according to claim 2, further comprising a table for matching each of transfer destinations stored therein with the unit for data division and a retransmission count, said packet reception information within said acknowledgement packet being referenced in order to write to the appropriate retransmission count field of said table the number of times data is not normally received by any specific destination, wherein, if any one of the stored retransmission counts exceeds a predetermined value, the corresponding data division unit is reduced.
5. A method for transmitting data via a network according to claim 3, further comprising a table for matching each of transfer destinations stored therein with the unit for data division and a transmission time, said table having a transmission time and a network load written therein every time a packet is transmitted, wherein, if any one of the stored transmission times exceeds a predetermined value, the corresponding data division unit is reduced.
6. A method for transmitting data via a network according to claim 3, wherein said transmitting side, having divided a data packet into a plurality of data packets to be transmitted, writes to each of the divided data packets the transmission time at which the packet in question is transmitted, wherein said receiving side adds said transmission time to an acknowledgement packet returned corresponding to the received data packet in question, and wherein said receiving side receives said acknowledgement packet and calculates the time required to transmit the packet by obtaining the difference between said transmission time and the current time.
7. A method for transmitting data via a network

according to claim 6, further comprising a table for matching each of transfer destinations stored therein with the unit for data division, a transmission time and a network load, said table having a transmission time and a network load written therein every time a packet is transmitted, wherein, if transmitting a packet by use of a large data division unit has made the corresponding network load in said table less than the current network load while increasing the corresponding transmission time, the corresponding data division unit in said table is replaced by a larger division unit.

8. A method for transmitting data via a network according to claim 1, wherein, when said network comprises two networks having different communication paths, said divided data packets are transferred via one network and said acknowledgement packet is transferred via the other network.
9. A method for transmitting data via a network according to claim 1, wherein said data packet is written to a nonvolatile area, wherein, before the transmission of said data packet is completed, a transmission end notice is sent to the source which requested the transmission of said data packet, and wherein, if the transmission of said data packet has failed, said data packet is retransmitted by use of the data held in said nonvolatile area.
10. A method for transmitting data via a network wherein data is transferred between a transmitting side (205(Fig. 2); 1903(Fig.22); 2704(Fig.27)) and a receiving side (206(Fig. 2); 1907(Fig.22); 2719(Fig.27));

said transmitting side (205,1903,2704) dividing one set of data into a plurality of sub-data, constructing a plurality of data packets each including one of said plurality of sub-data, and sending said data packets to said transmitting side;

said receiving side (206,1907,2719) receiving each of the transmitted packets and returning to said transmitting side an acknowledgement including a result of reception regarding each transmitted data packet;

said transmitting side (205,1903,2704) receiving said acknowledgement and retransmitting a data packet corresponding to a received acknowledgement which indicates failure of transmission; and

said receiving side (206,1903,2704) reconstructing said one set of data from said sub-data in said plurality of data packets, after having normally received all packets constituting said one set of data.

FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5
PRIOR ART




FIG. 6
PRIOR ART




FIG. 7




FIG. 8

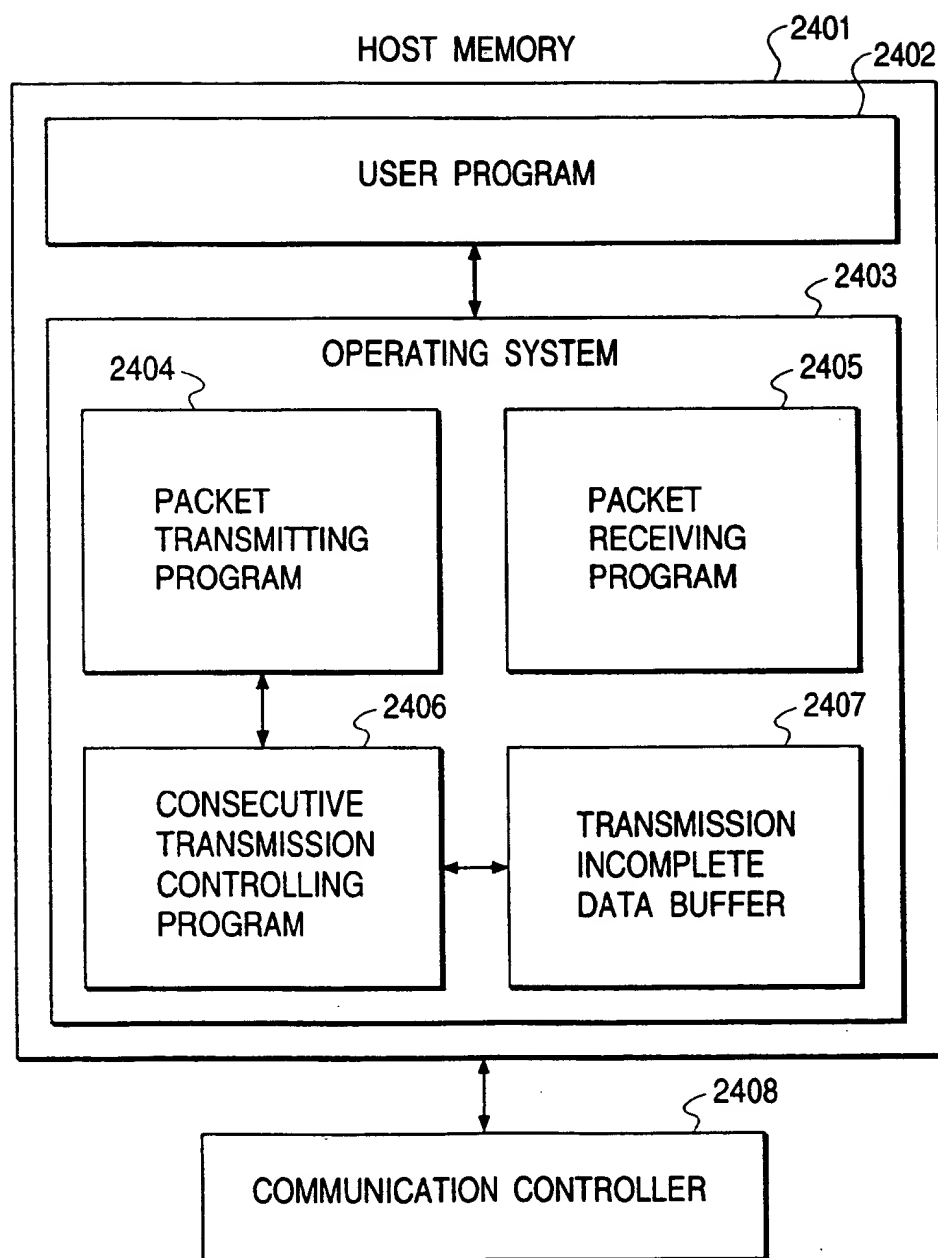


FIG. 9

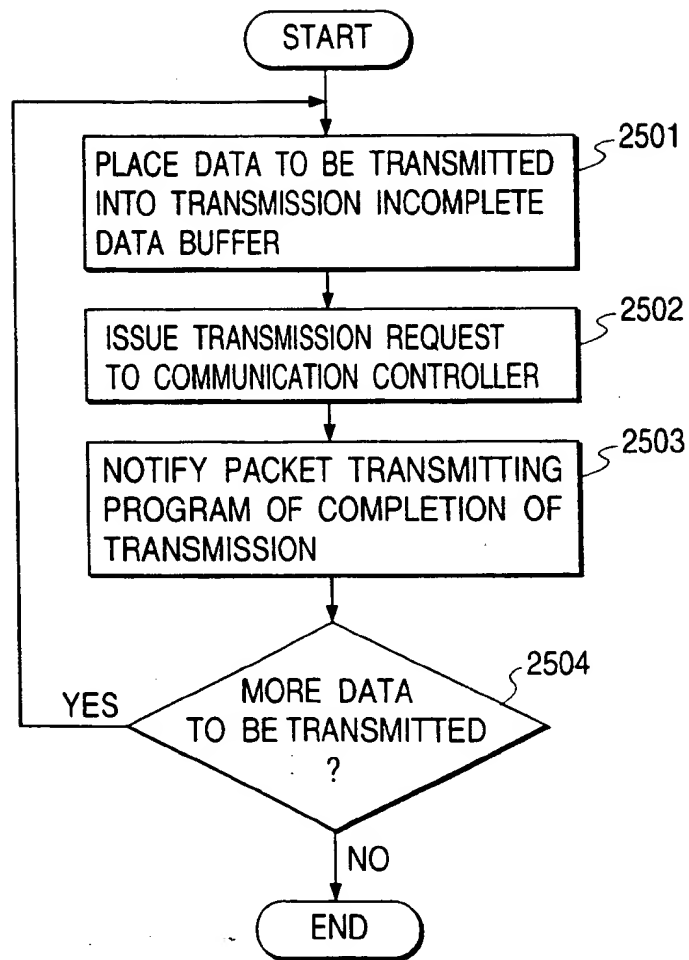


FIG. 10

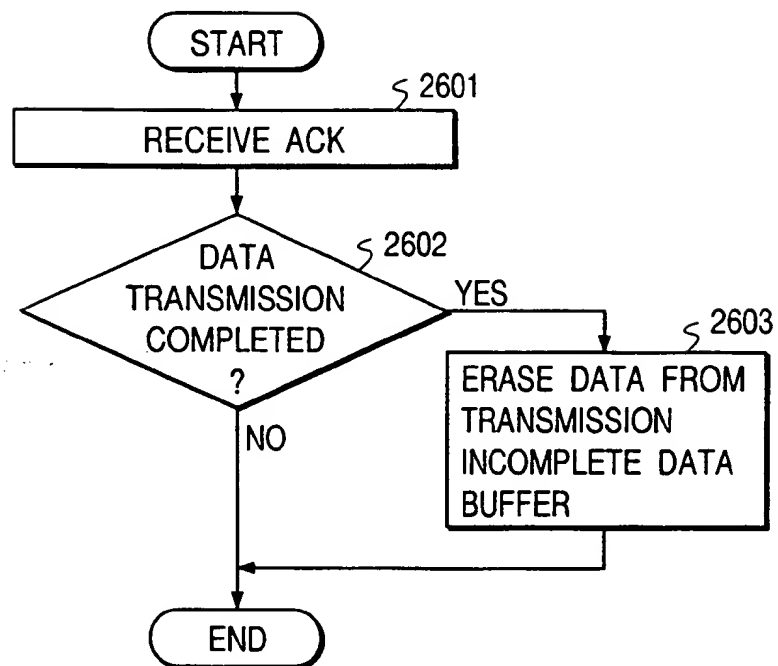


FIG. 11




FIG. 12

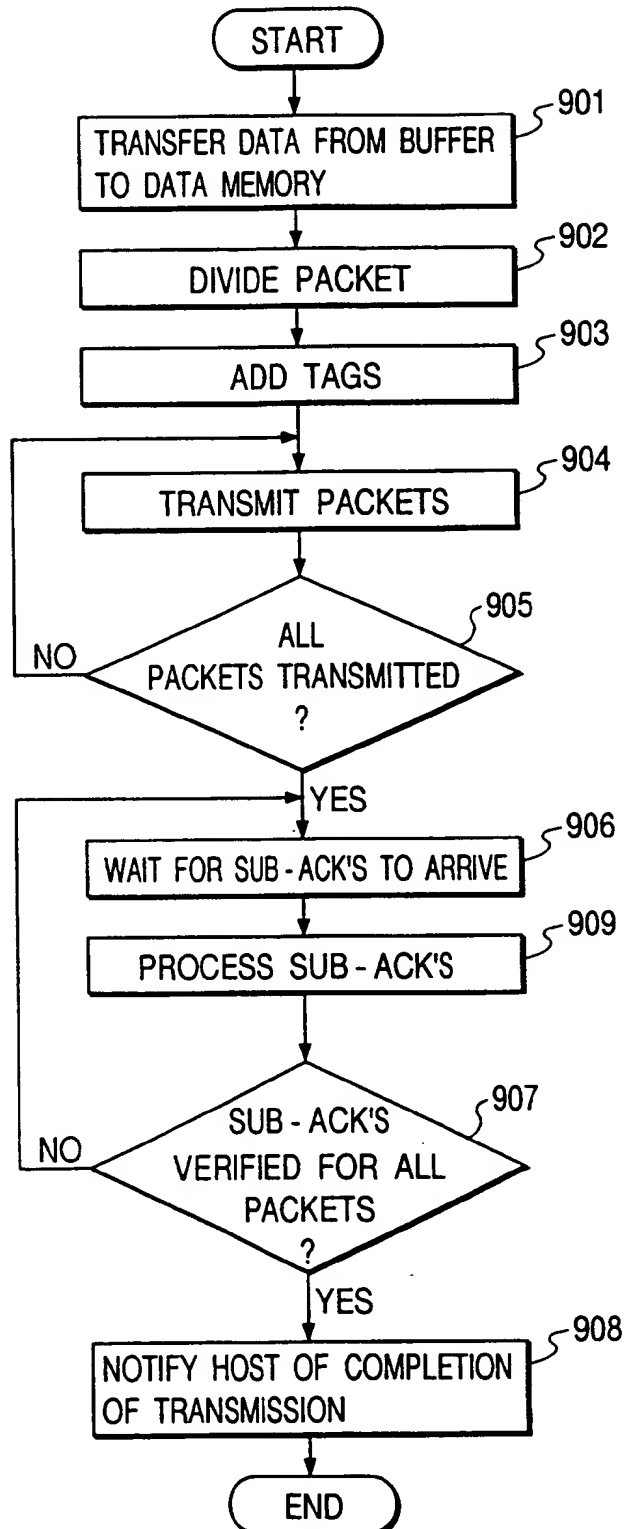


FIG. 13

DESTINATION ID	DIVISION UNIT	RESPONSE TIME	COMMUNI - CATION LOAD	DIVISION DIRECTION
A	2KB	10ms	10%	+1
B	4KB	20ms	30%	+1
C	16KB	50ms	10%	-1
-	N	N	N	N
-	N	N	N	N

FIG. 14

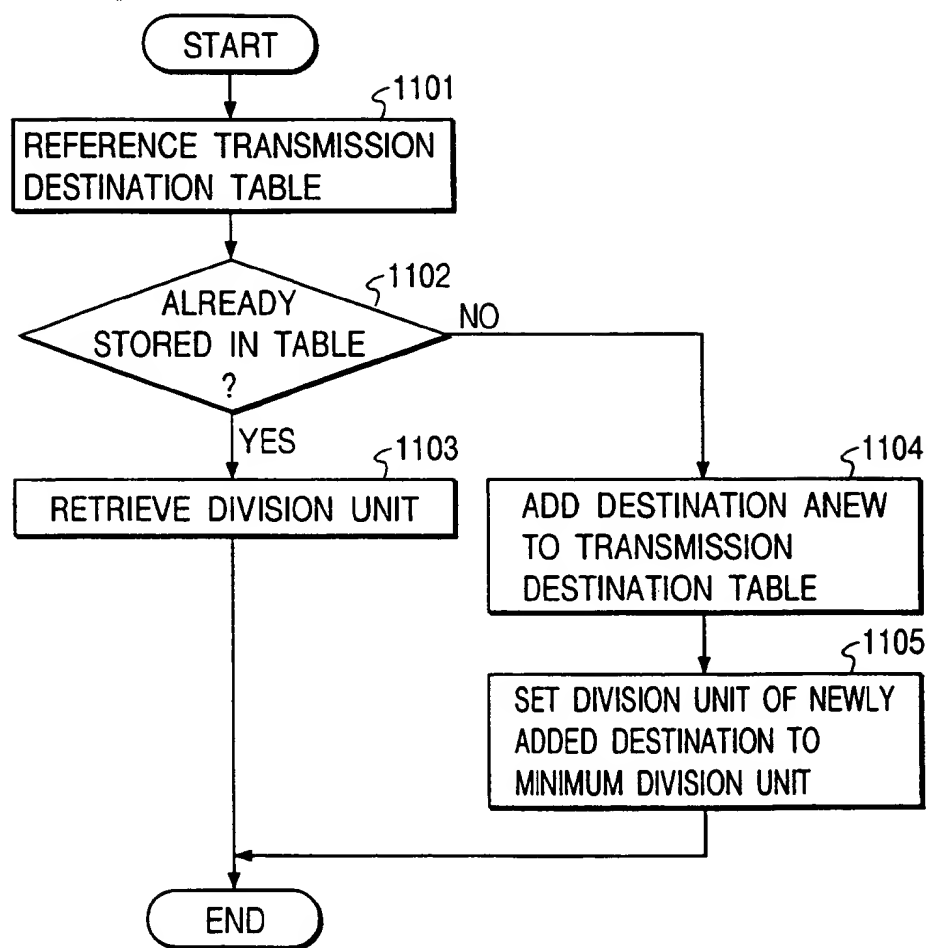


FIG. 15

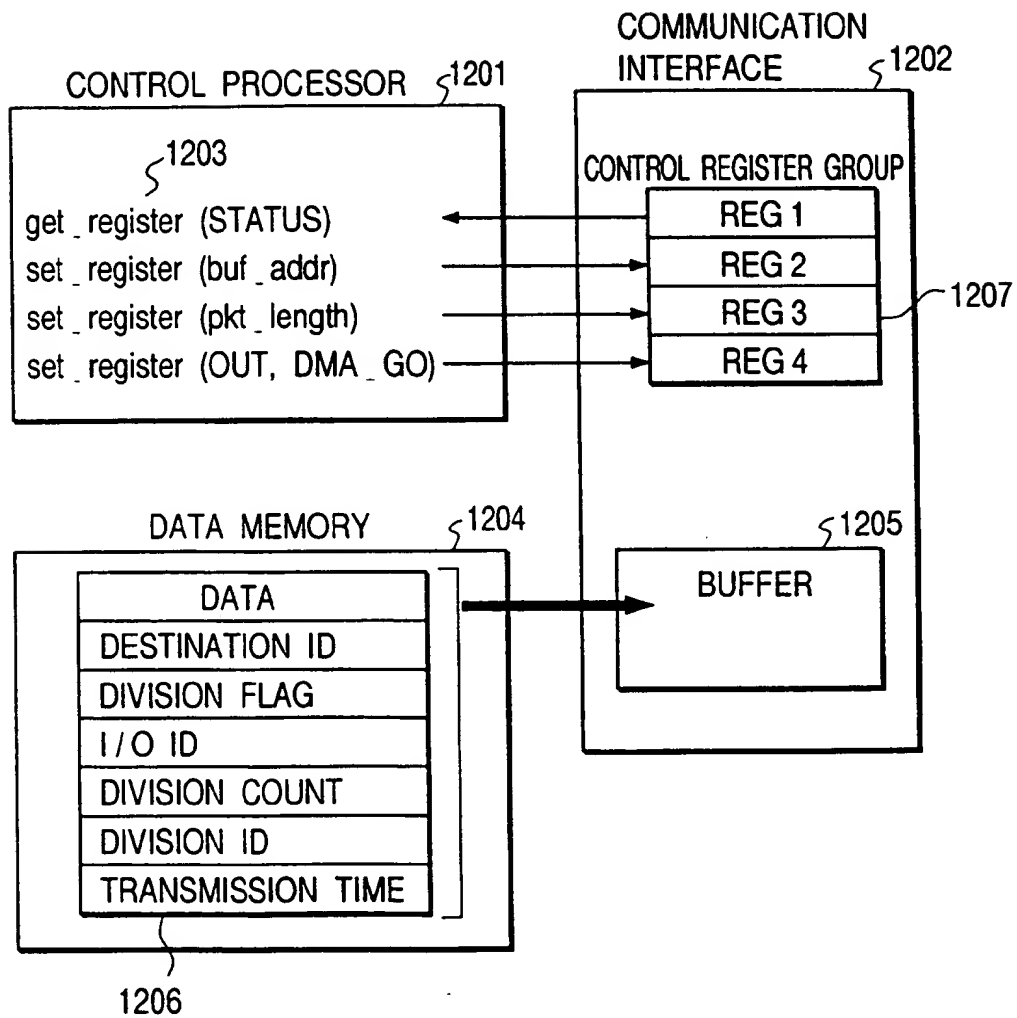


FIG. 16

RESULT OF TRANSMISSION	1301
DESTINATION ID	1302
DIVISION FLAG	1303
PACKET ID	1304
DIVISION COUNT	1305
DIVISION ID	1306
TRANSMISSION TIME	1307

FIG. 17

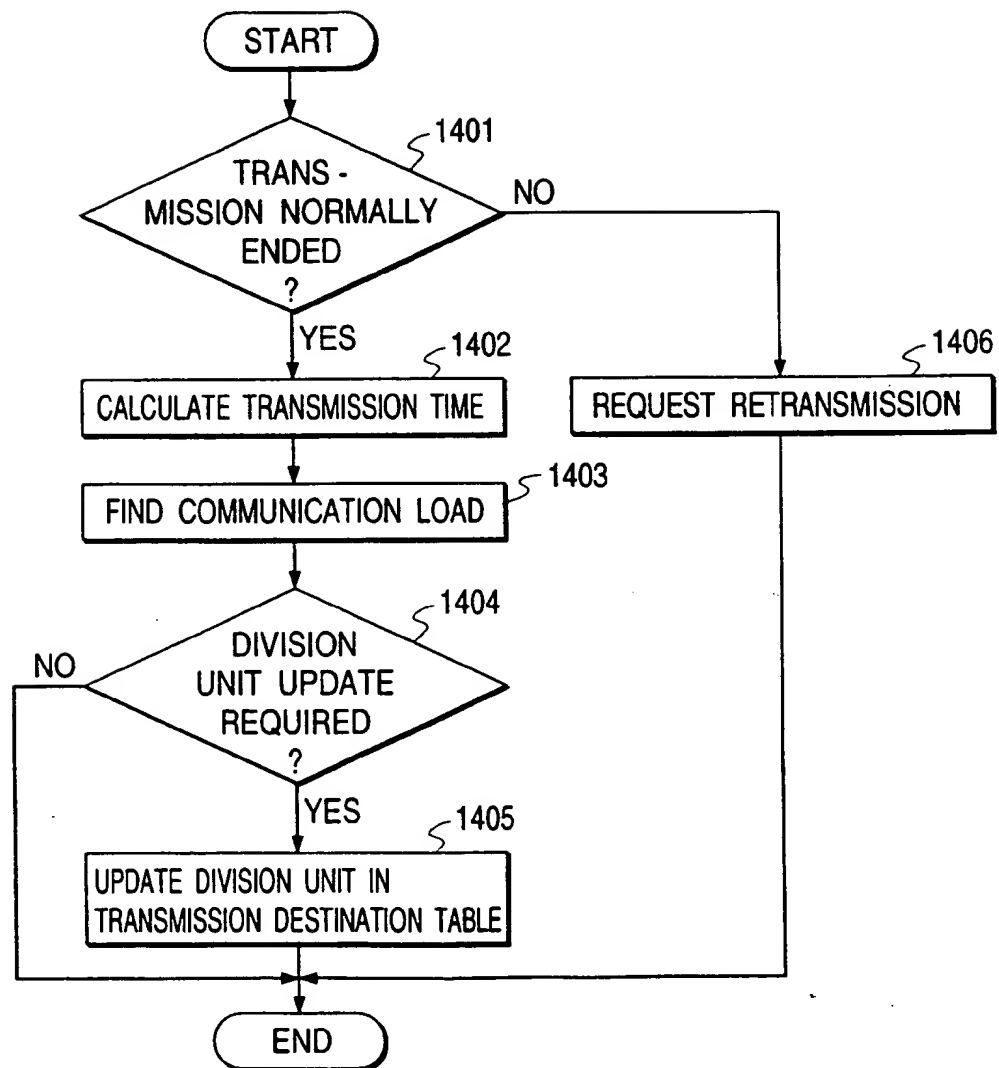


FIG. 18

PACKET ID	DIVISION ID	ACK RECEIPT	RETRANSMISSION COUNT
1	1	RECEIVED	0
	2	RECEIVED	0
	3	YET TO BE RECEIVED	2
	4	RECEIVED	3
2	1	YET TO BE RECEIVED	0
	2	RECEIVED	0
	3	RECEIVED	0
...	

FIG. 19




FIG. 20

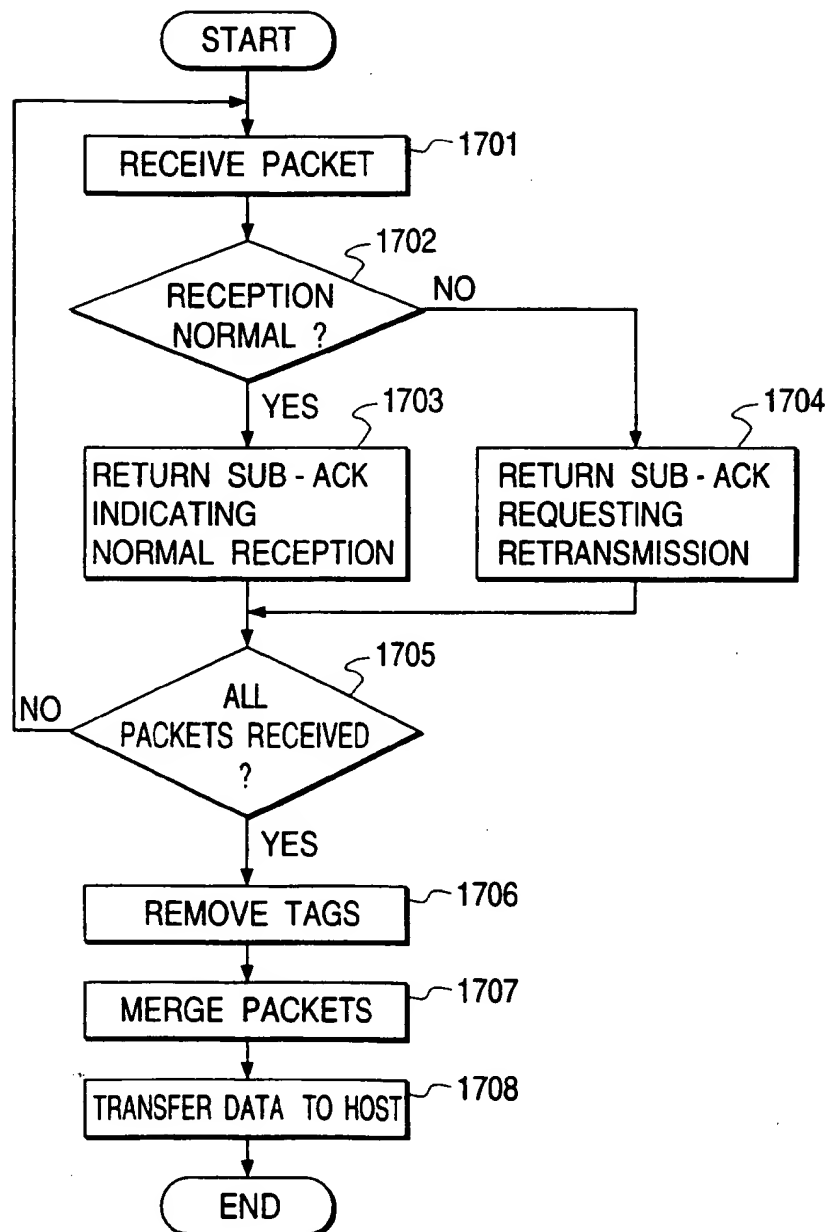


FIG. 21

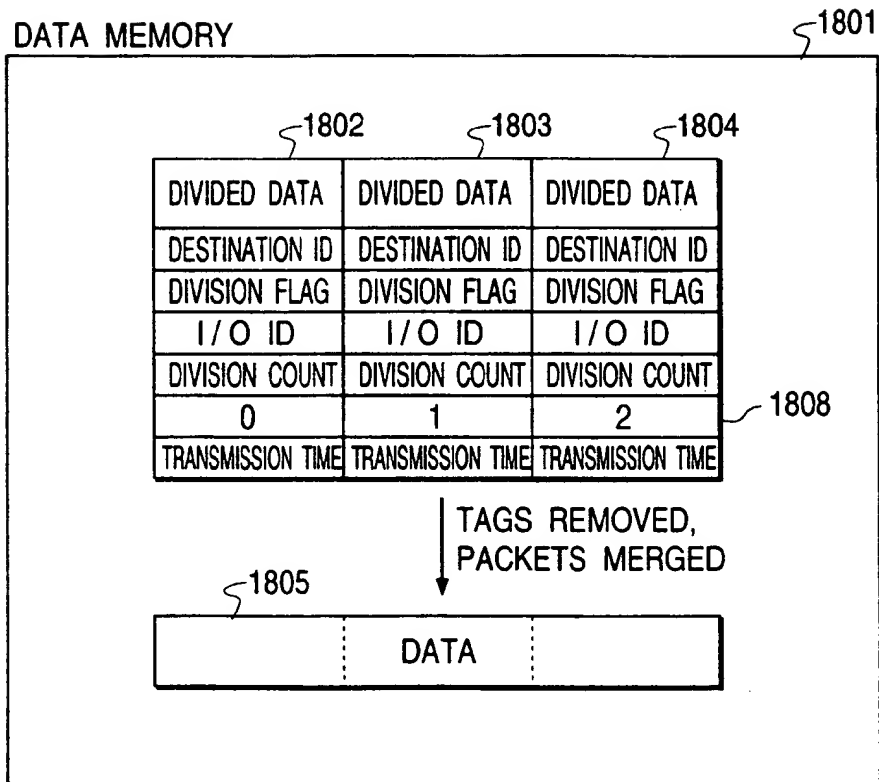


FIG. 22




FIG. 23

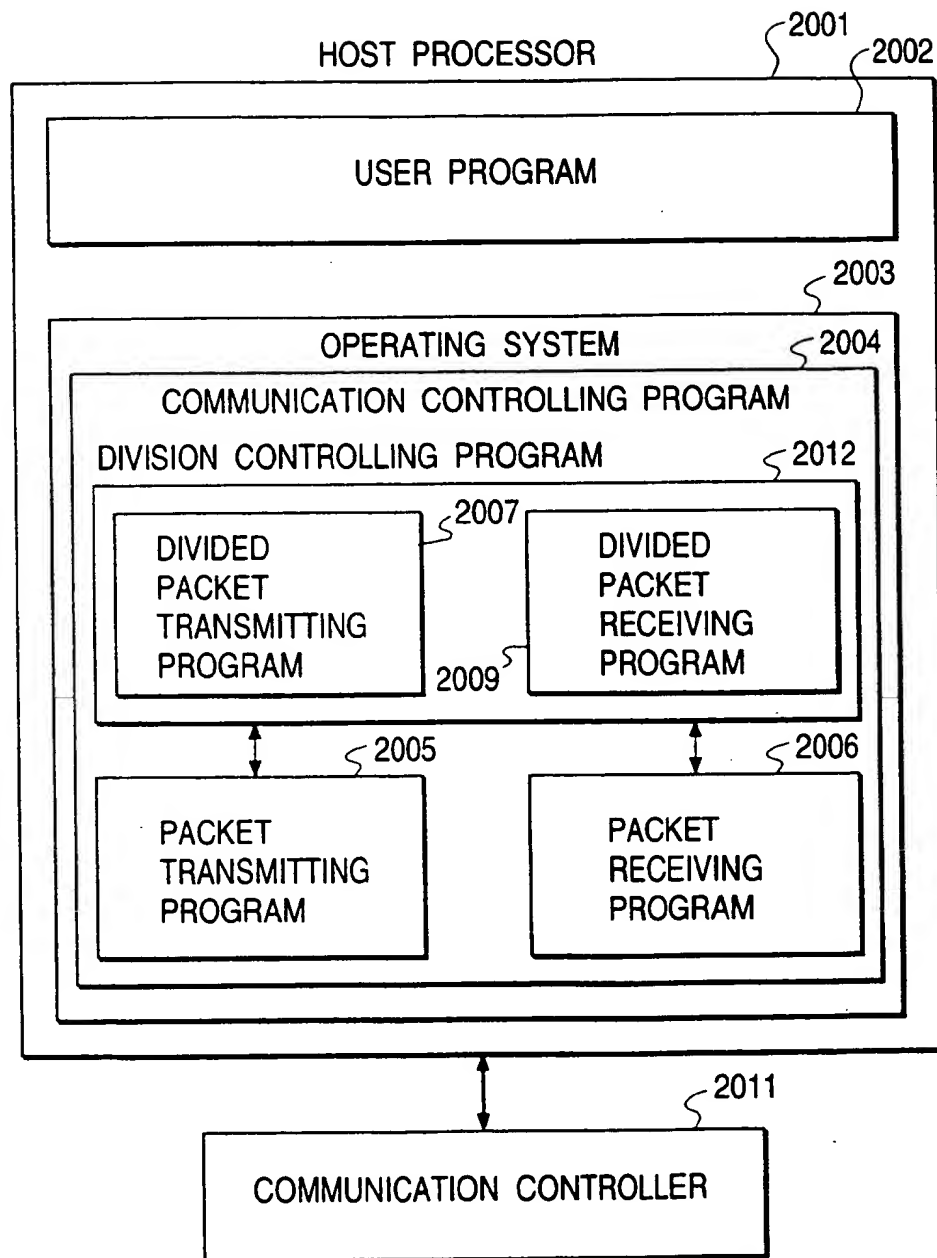


FIG. 24




FIG. 25

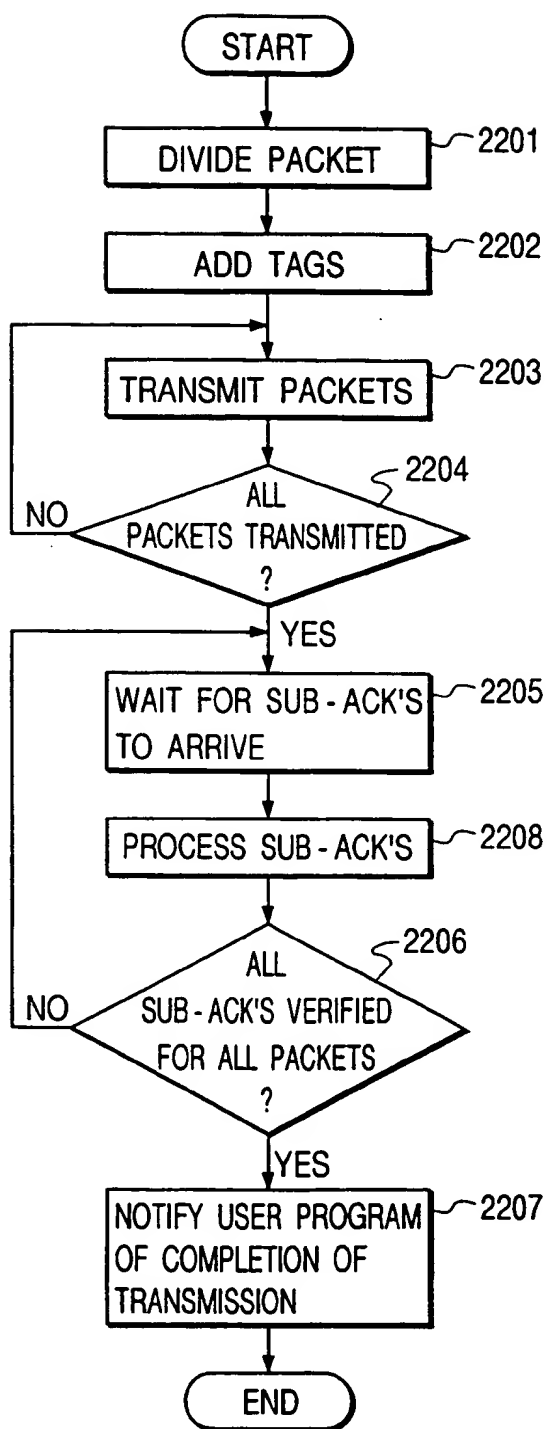


FIG. 26

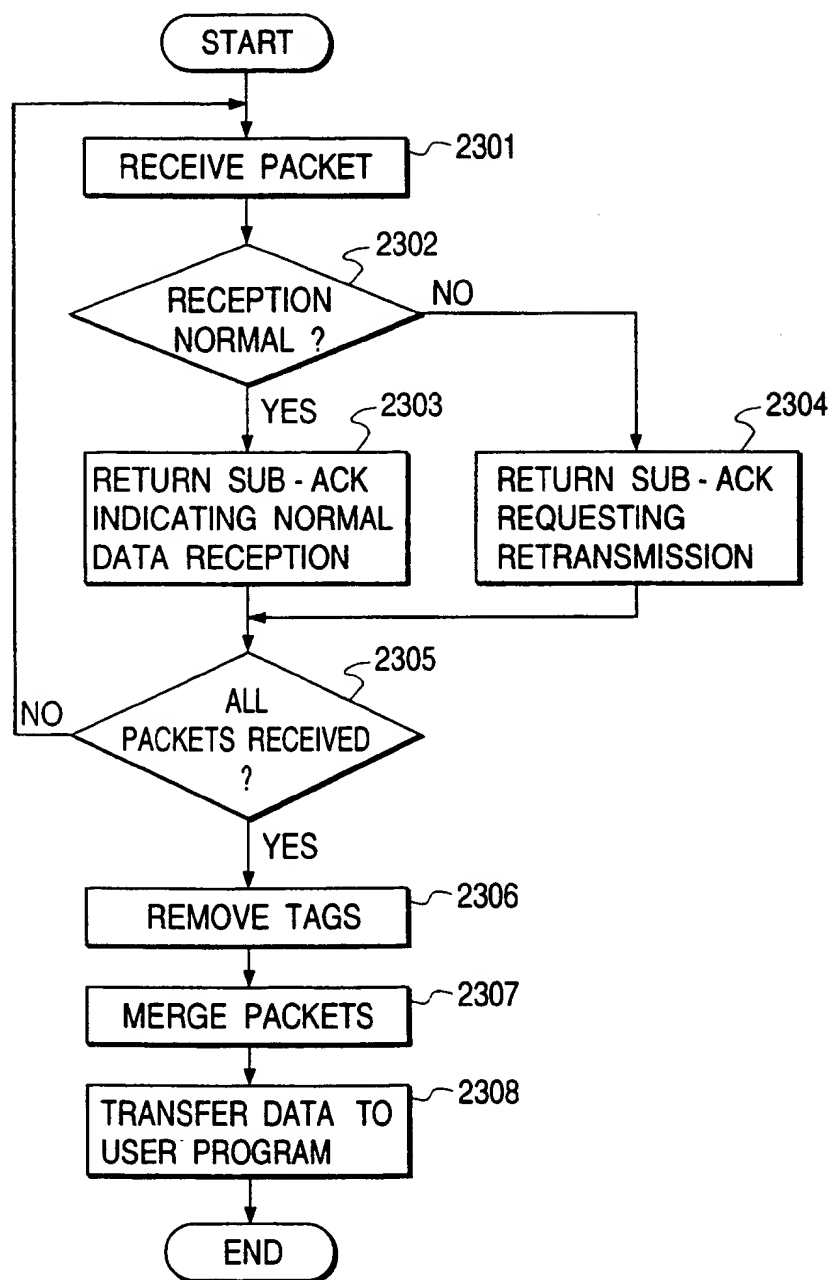


FIG. 27




FIG. 1

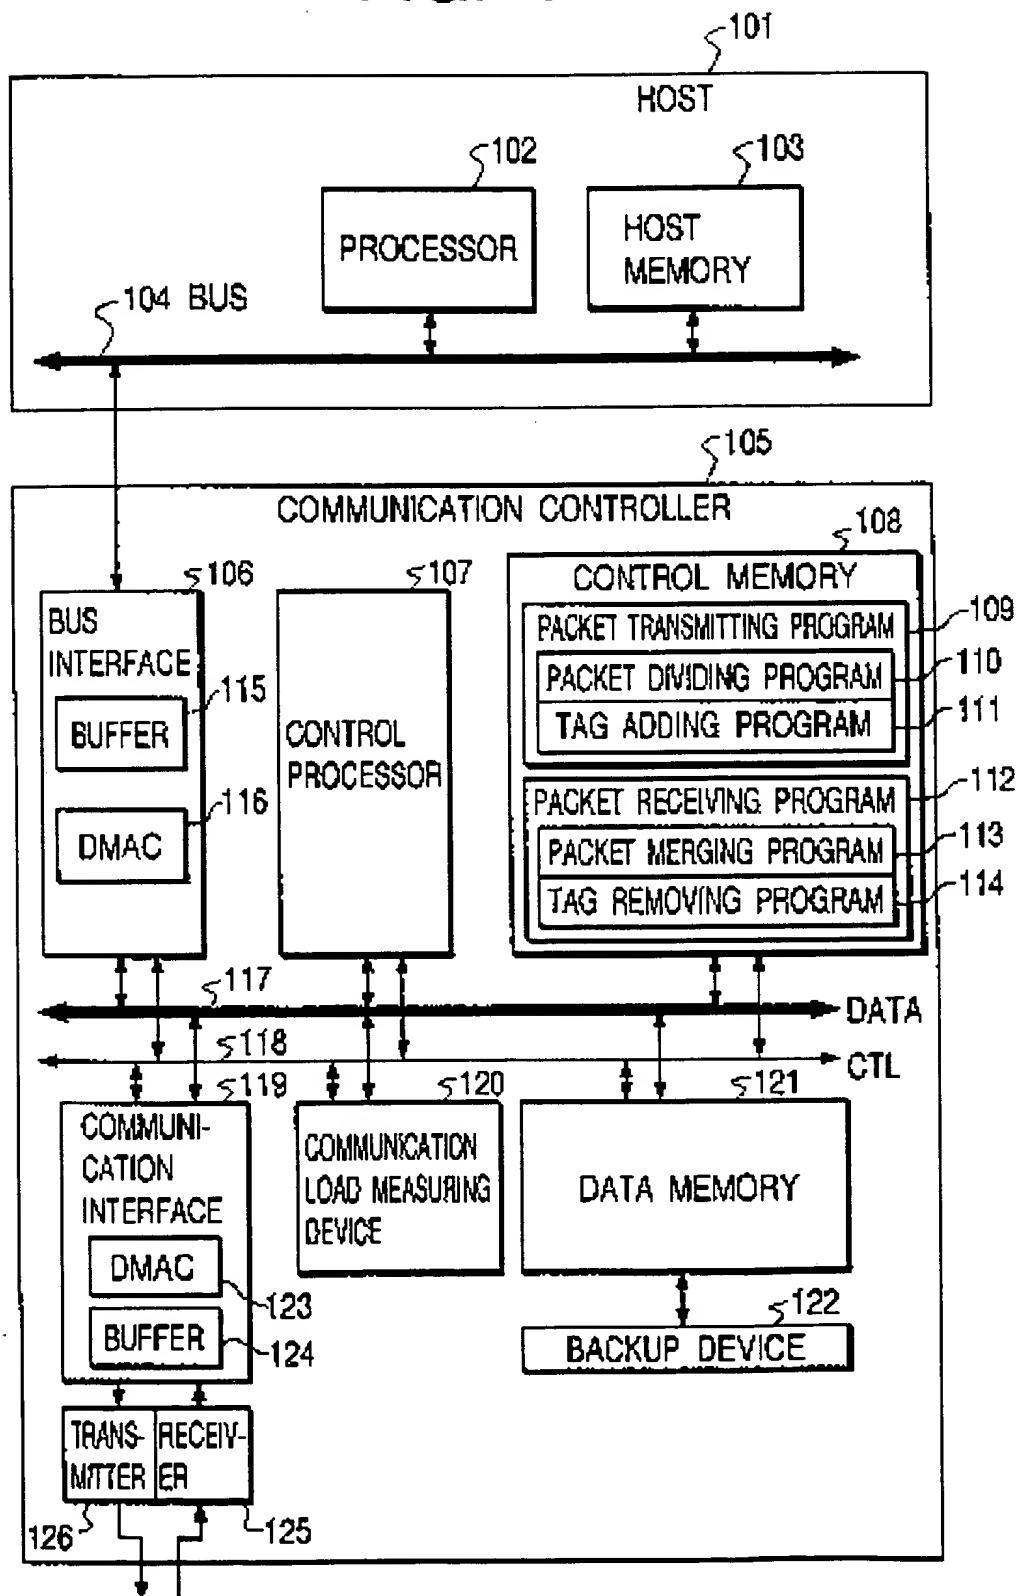


FIG. 2




FIG. 3




FIG. 4




FIG. 5
PRIOR ART




FIG. 6
PRIOR ART




FIG. 7

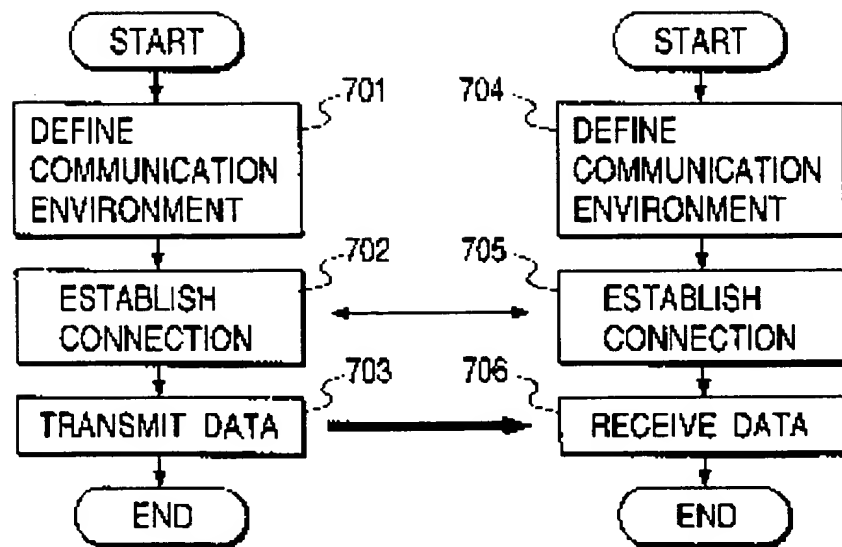


FIG. 8

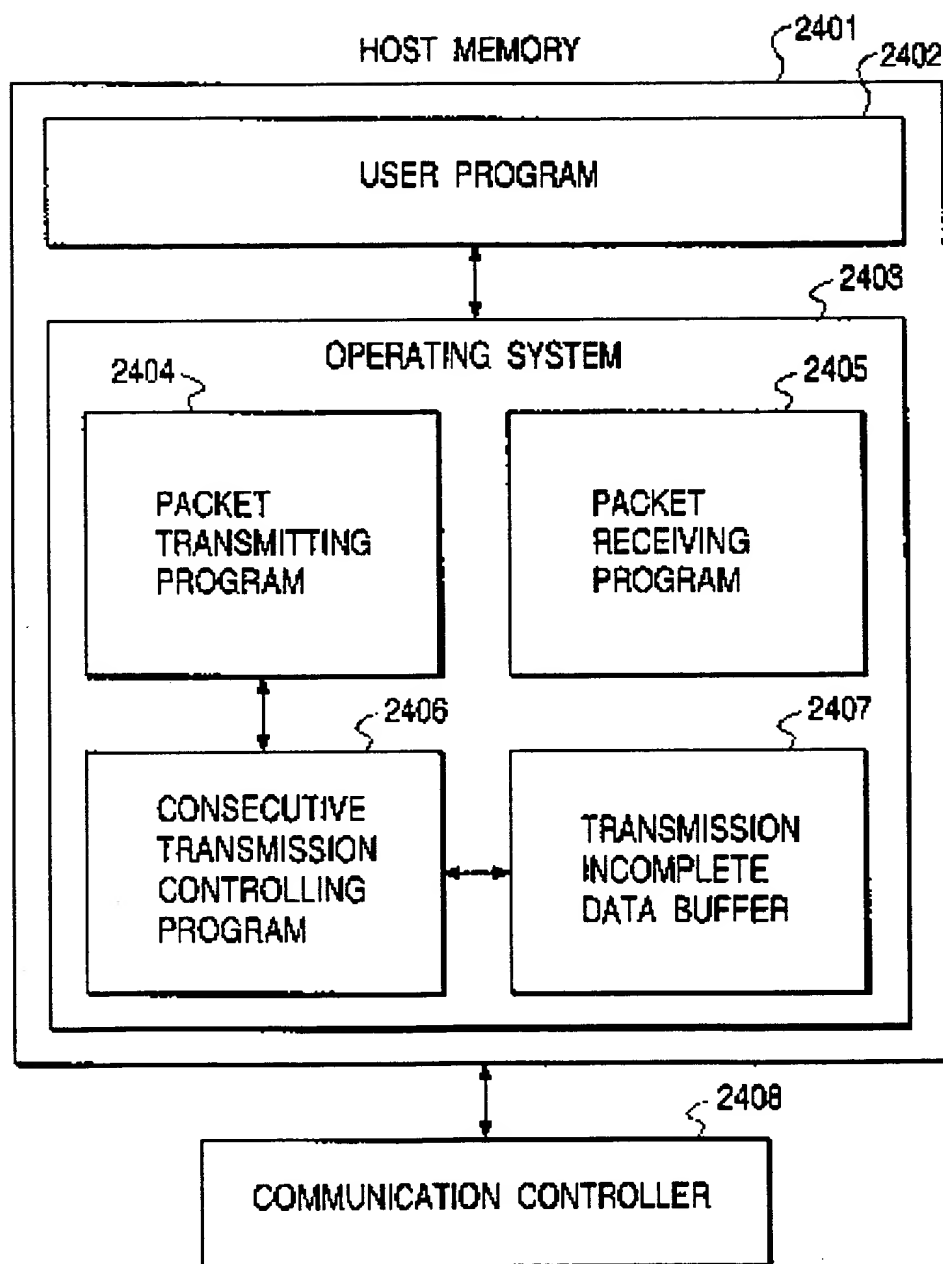


FIG. 9

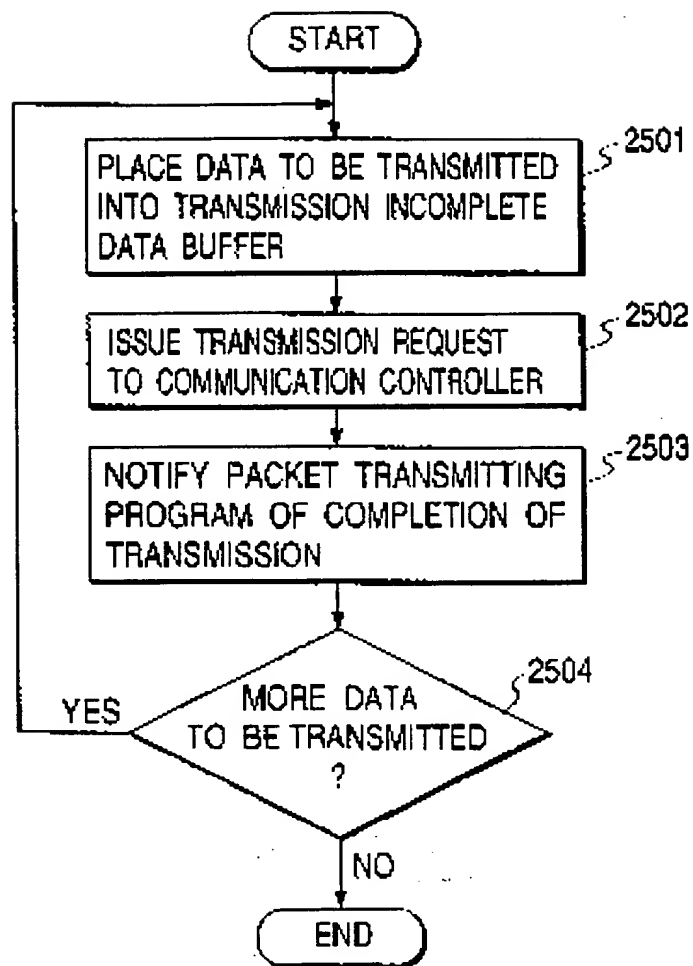


FIG. 10

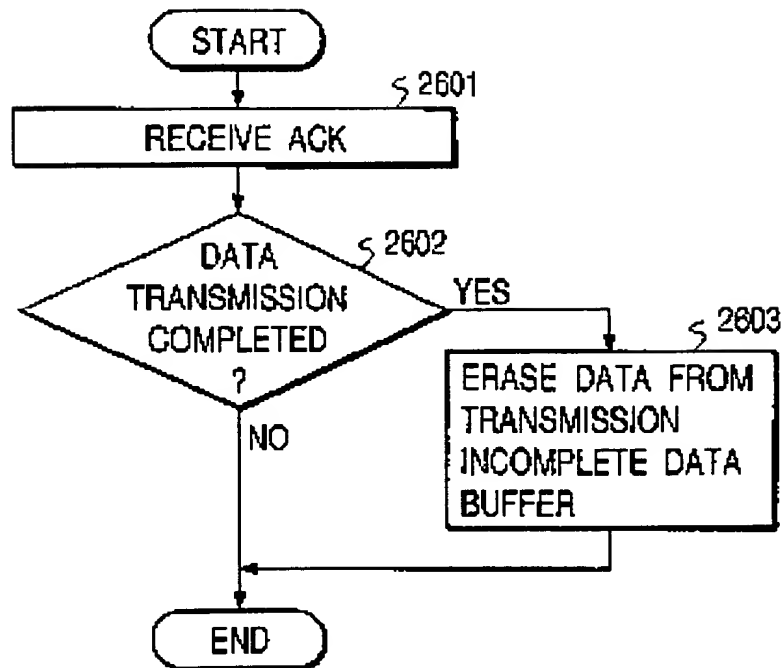


FIG. 11




FIG. 12




FIG. 13

DESTINATION ID	DIVISION UNIT	RESPONSE TIME	COMMUNI - CATION LOAD	DIVISION DIRECTION
A	2KB	10ms	10%	+1
B	4KB	20ms	30%	+1
C	16KB	50ms	10%	-1
.	N	N	N	N
.	N	N	N	N

FIG. 14

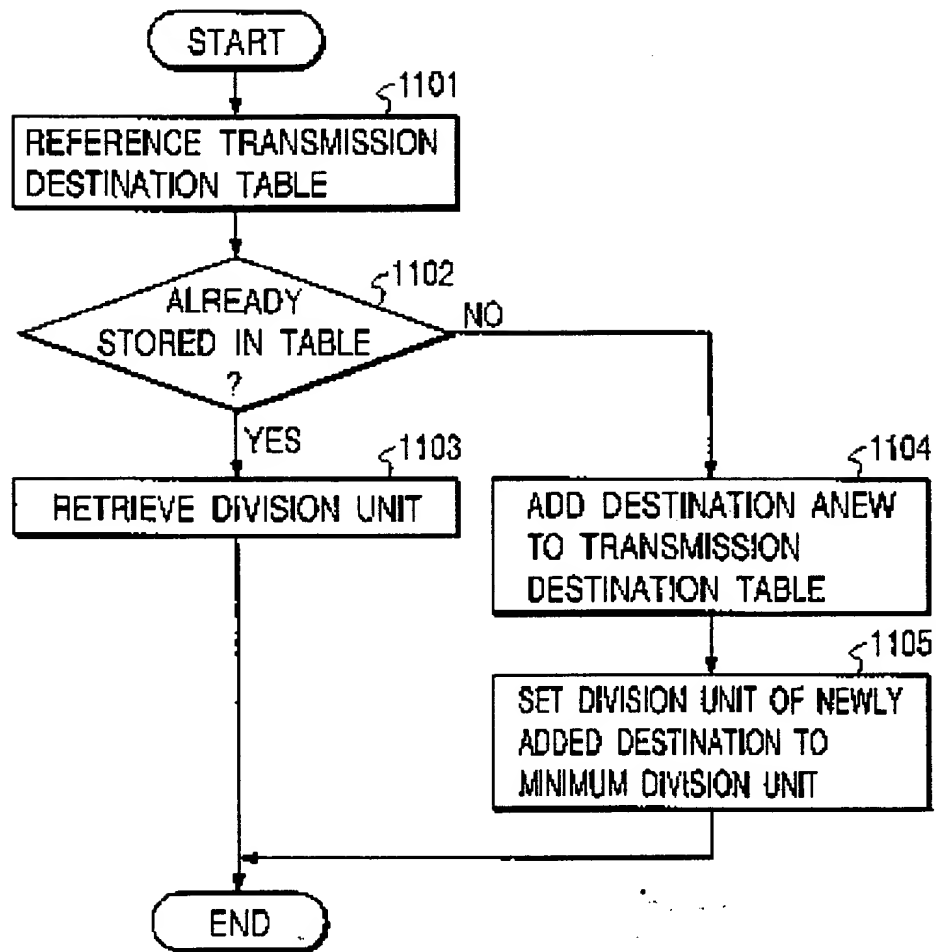


FIG. 15

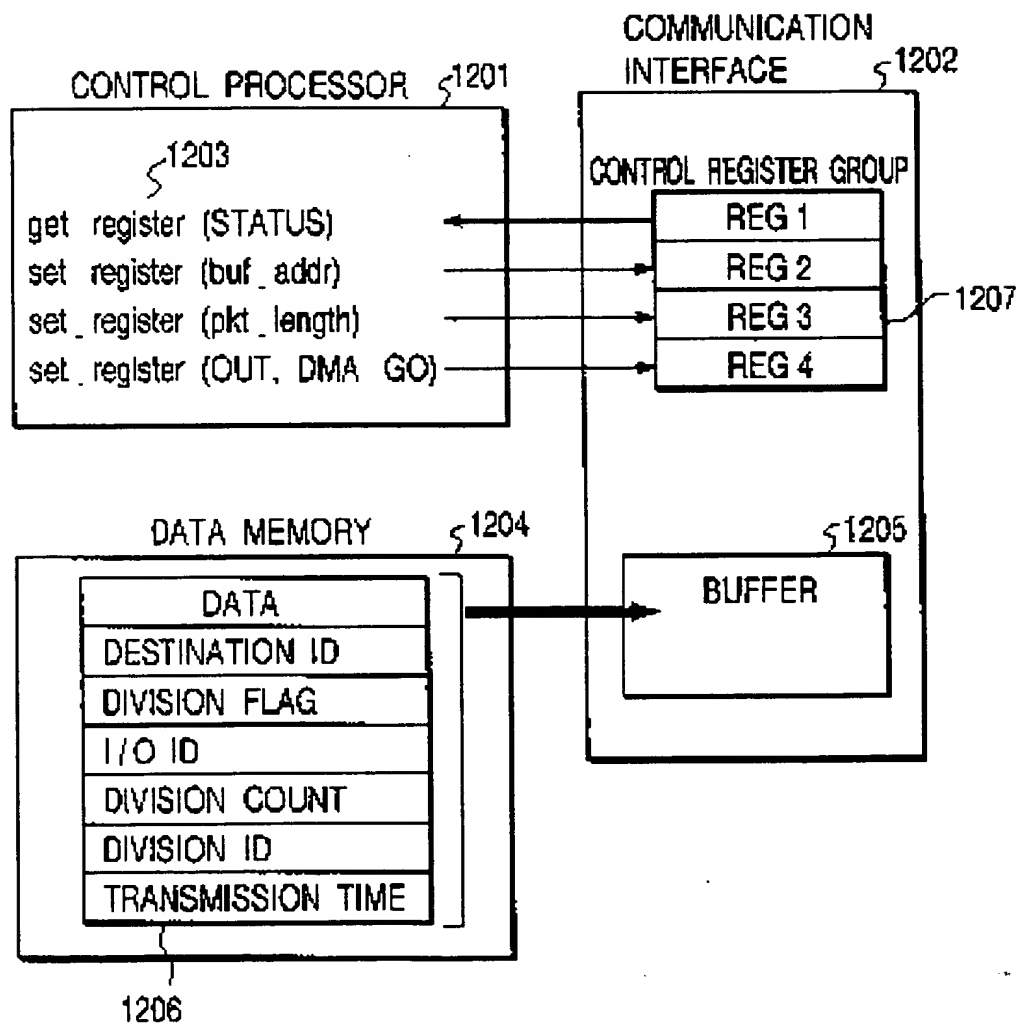


FIG. 16

RESULT OF TRANSMISSION	1301
DESTINATION ID	1302
DIVISION FLAG	1303
PACKET ID	1304
DIVISION COUNT	1305
DIVISION ID	1306
TRANSMISSION TIME	1307

FIG. 17

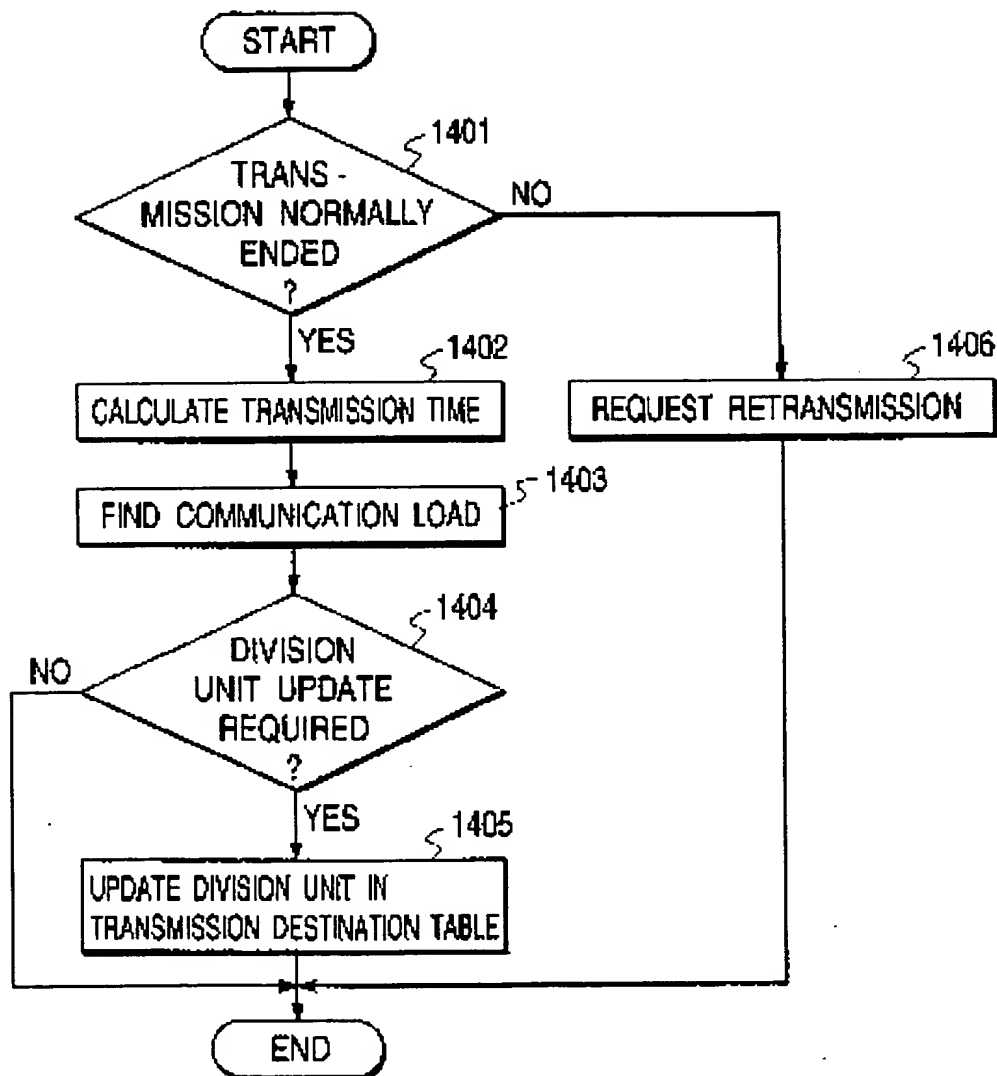


FIG. 18

PACKET ID	DIVISION ID	ACK RECEIPT	RETRANSMISSION COUNT
1	1	RECEIVED	0
	2	RECEIVED	0
	3	YET TO BE RECEIVED	2
	4	RECEIVED	3
2	1	YET TO BE RECEIVED	0
	2	RECEIVED	0
	3	RECEIVED	0
...	

FIG. 19




FIG. 20

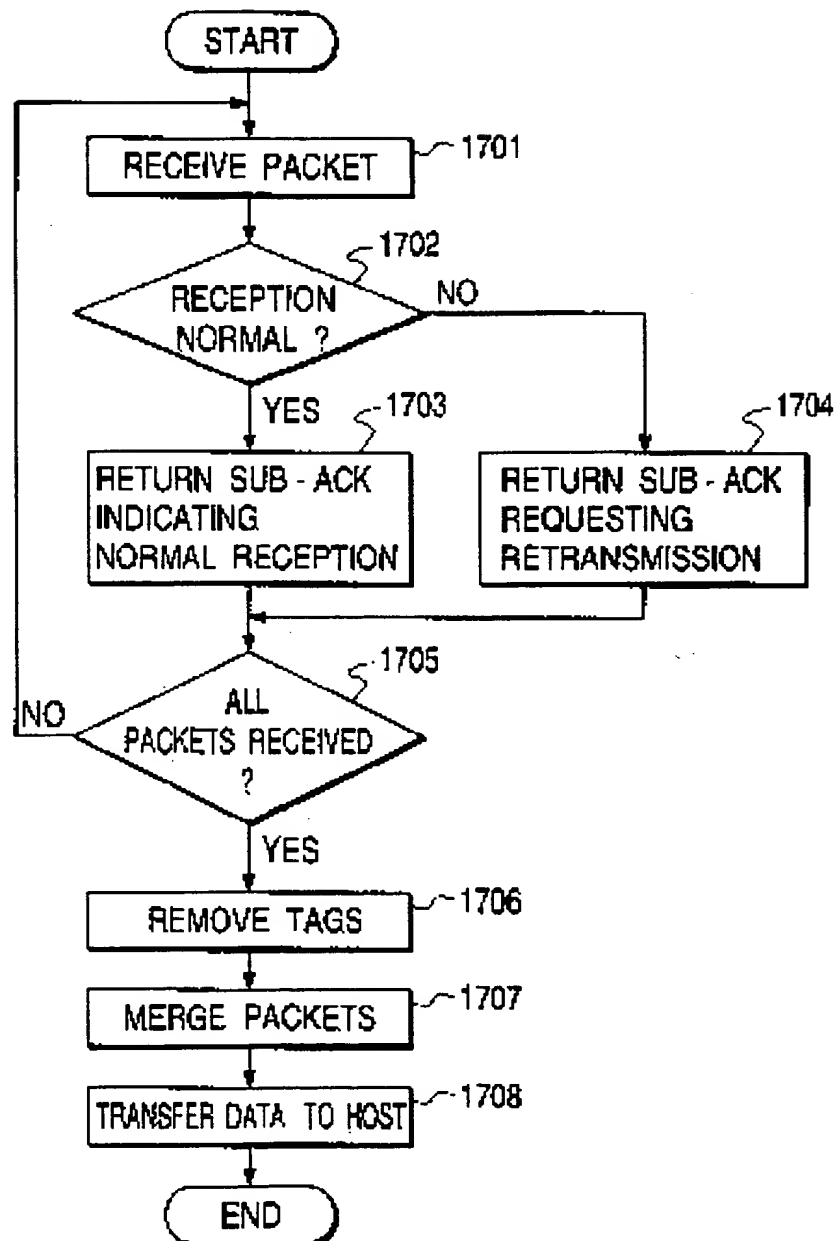


FIG. 21




FIG. 22




FIG. 23

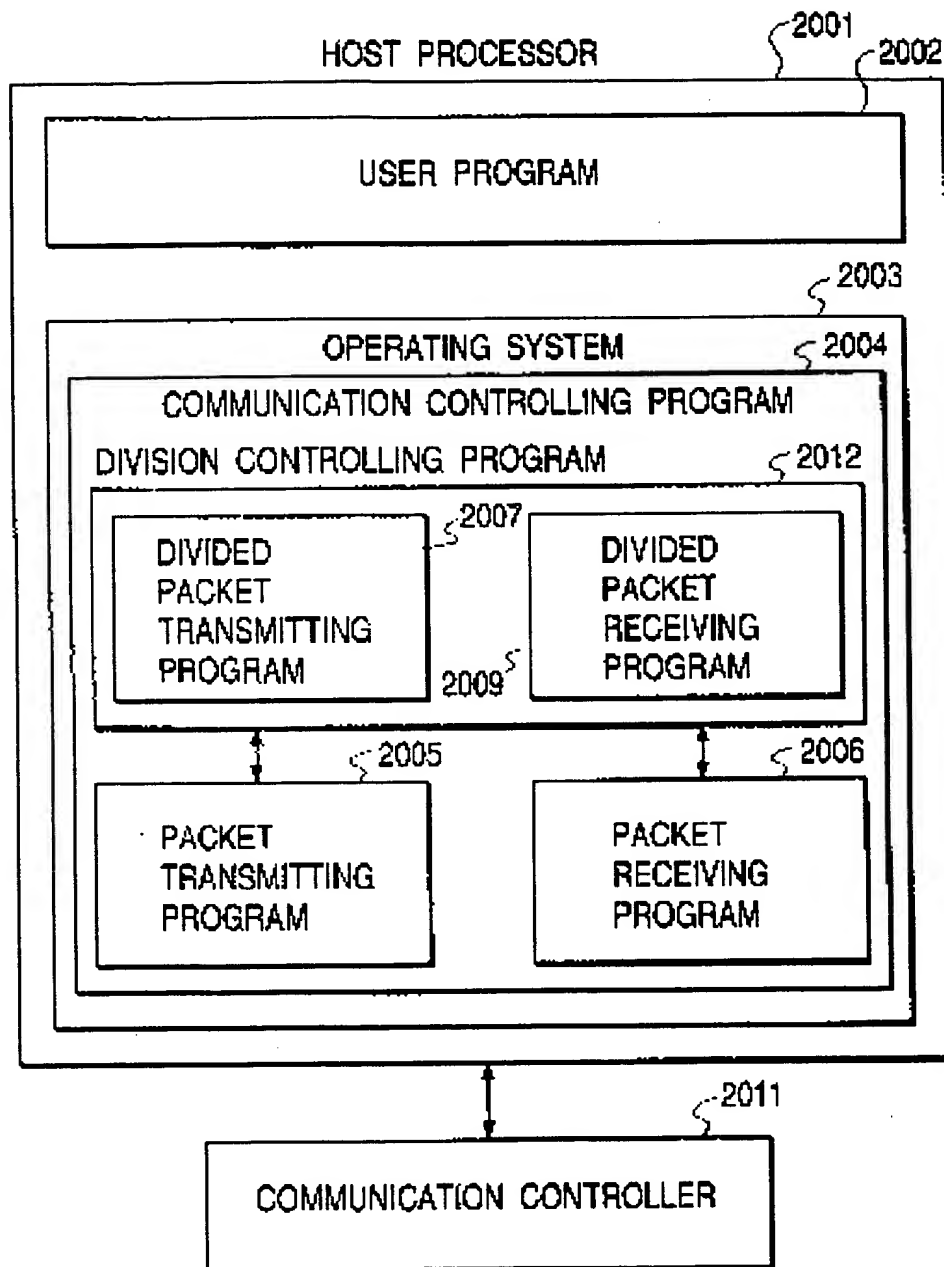


FIG. 24




FIG. 25

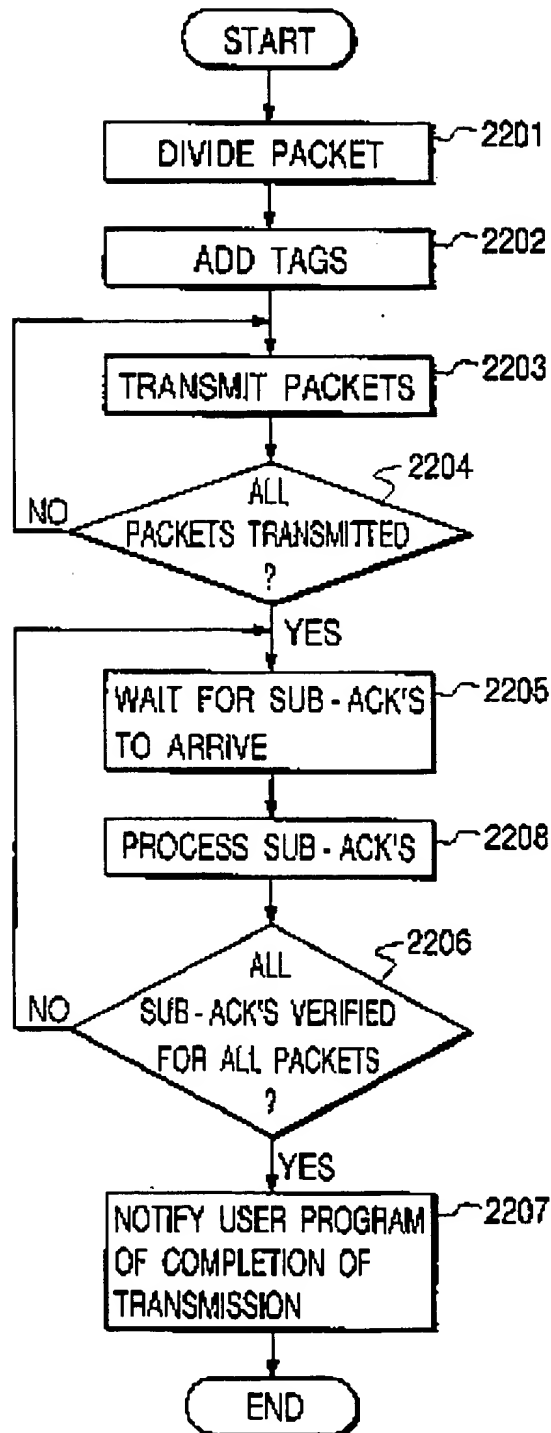


FIG. 26






FIG. 27





(19)

Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 0 762 705 A3

(12)

EUROPEAN PATENT APPLICATION

(88) Date of publication A3:
27.09.2000 Bulletin 2000/39

(51) Int. Cl.⁷: H04L 29/06, H04L 1/18

(43) Date of publication A2:
12.03.1997 Bulletin 1997/11

(21) Application number: 96114234.6

(22) Date of filing: 05.09.1996

(84) Designated Contracting States:
DE FR GB LU MC NL PT SE

(30) Priority: 08.09.1995 JP 23102395

(71) Applicant: HITACHI, LTD.
Chiyoda-ku, Tokyo (JP)

(72) Inventors:
• Takamoto, Yoshifumi
Fuchu-Shi, Tokyo (JP)
• Kanai, Hiroki
Higashi-Yamato-Shi, Tokyo (JP)


• Takase, Tadahiyo
Oita-Ken, Tokyo (JP)
• Kitai, Katsuyoshi
Kokubunji-Shi, Tokyo (JP)
• Masuoka, Yoshimasa
Kodaira-Shi, Tokyo (JP)

(74) Representative:
Beetz & Partner
Patentanwälte
Steinsdorfstrasse 10
80538 München (DE)

(54) Method for transmitting data via a network

(57) A communication controller (205) at the sending computer (201) divides data transferred from a host (101) into sub-ACK unit packets, and transfers them sequentially to a destination without waiting for the sub-ACK's being subsequently provided by the destination. A communication controller (206) at the receiving computer (202) issues a sub-ACK to the sending computer (201) for each of the sub-ACK unit packet, if the sub-ACK packet has been normally received and otherwise issues retransmission request for the sub-ACK unit packets and merges data included in the sub-ACK unit packets into the initial data, after they are normally received.

FIG. 2



EP 0 762 705 A3



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 96 11 4234

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (IntCl.6)
X	US 5 440 545 A (HAMILTON JR R LEE ET AL) 8 August 1995 (1995-08-08) * the whole document *	1-4,9,10	H04L29/06 H04L1/18
A	---	8	
A	US 5 396 490 A (WHITE RICHARD E ET AL) 7 March 1995 (1995-03-07) * column 5, line 12 - column 6, line 48 * * abstract; figures 2-6,10 *	1-4,9,10	
A	--- "METHODOLOGY TO EXECUTE AND FORWARD APPLICATION COMMANDS AT REMOTE WORKSTATIONS" March 1992 (1992-03) , IBM TECHNICAL DISCLOSURE BULLETIN,US,IBM CORP. NEW YORK, VOL. 34, NR. 10A, PAGE(S) 44-46 XP000302216 ISSN: 0018-8689 * the whole document *	5-7	
A	US 4 484 326 A (TURNER JONATHAN S) 20 November 1984 (1984-11-20) * column 2, line 43 - column 11, line 2 * * figures 3,4,6,7 *	5-7	TECHNICAL FIELDS SEARCHED (Int.Cl.6) H04L
A	--- PATENT ABSTRACTS OF JAPAN vol. 012, no. 124 (E-601), 16 April 1988 (1988-04-16) & JP 62 248337 A (NIPPON TELEGR & TELEPH CORP), 29 October 1987 (1987-10-29) * abstract *	1-10	

		-/--	
The present search report has been drawn up for all claims			
Place of search BERLIN		Date of completion of the search 31 July 2000	Examiner Kalabic, F
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document</p>			

EPO FORM 1503 03 82 (P04C01)



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 96 11 4234

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (IntCl.6)
A	<p>"Protocol for Isochronous Traffic over Fiber Channel Switching"</p> <p>IBM TECHNICAL DISCLOSURE BULLETIN, vol. 37, no. 6b, June 1994 (1994-06), pages 377-380, XP002143492</p> <p>IBM corp. New York</p> <p>* the whole document *</p> <p>-----</p>	5-7	
			TECHNICAL FIELDS SEARCHED (IntCl.6)
The present search report has been drawn up for all claims			
Place of search		Date of completion of the search	Examiner
BERLIN		31 July 2000	Kalabic, F
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone</p> <p>Y : particularly relevant if combined with another document of the same category</p> <p>A : technological background</p> <p>O : non-written disclosure</p> <p>P : intermediate document</p> <p>T : theory or principle underlying the invention</p> <p>E : earlier patent document, but published on, or after the filing date</p> <p>D : document cited in the application</p> <p>L : document cited for other reasons</p> <p>& : member of the same patent family, corresponding document</p>			

EPO FORM 1503 03.82 (P04001)

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 96 11 4234

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.
The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

31-07-2000

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5440545 A	08-08-1995	AU 675868 B	20-02-1997
		AU 7321094 A	28-02-1995
		BR 9405540 A	08-09-1999
		CA 2143950 A	09-02-1995
		CN 1113402 A,B	13-12-1995
		EP 0667074 A	16-08-1995
		JP 8502159 T	05-03-1996
		KR 197922 B	15-06-1999
		WO 9504421 A	09-02-1995
US 5396490 A	07-03-1995	BR 9306050 A	18-11-1997
		EP 0632945 A	11-01-1995
		JP 7505509 T	15-06-1995
		KR 159795 B	01-12-1998
		WO 9319544 A	30-09-1993
US 4484326 A	20-11-1984	CA 1208749 A	29-07-1986
		EP 0108554 A	16-05-1984
		EP 0124591 A	14-11-1984
		JP 2023109 B	22-05-1990
		JP 59501851 T	01-11-1984
		WO 8401869 A	10-05-1984
JP 62248337 A	29-10-1987	NONE	

EPO FORM P0458

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82